On Budget-Feasible Mechanism Design for Symmetric Submodular Objectives

Georgios Amanatidis Georgios Birmpas Evangelos Markakis
Athens University of Economics and Business, Department of Informatics.
{gamana, gebirbas, markakis}@aueb.gr

April 25, 2017

arXiv:1704.06901v1 [cs.GT] 23 Apr 2017

Abstract

We study a class of procurement auctions, where an auctioneer is interested in buying resources or services from a set of agents. The auctioneer is interested in selecting a subset of agents so as to maximize her valuation function, without exceeding a given budget. As the resources are owned by strategic agents, our overall goal is to design mechanisms that are i) truthful, ii) budget-feasible, and iii) obtain a good approximation to the optimal value. Budget-feasibility creates more challenges, making several approaches inapplicable in this setting.

Previous results on budget-feasible mechanisms have considered mostly monotone valuation functions. In this work, we mainly focus on *symmetric submodular* valuations, a prominent class of non-monotone submodular functions, that includes cut functions. We begin first with a purely algorithmic result, obtaining a $\frac{2e}{e-1}$-approximation for symmetric submodular functions under a budget constraint. We view this as a standalone result of independent interest, as it is the best known factor achieved by a deterministic algorithm. We then proceed to propose truthful, budget feasible mechanisms (both deterministic and randomized), paying particular attention on the Budgeted Max Cut problem. Our results significantly improve the known approximation ratios for these objectives, while establishing polynomial running time for many cases where only exponential mechanisms were known. For most symmetric submodular functions, only randomized mechanisms with very large approximation ratios were known prior to our results. At the heart of our approach lies an appropriate combination of local search algorithms with results for monotone submodular valuations, applied to the derived local optima.

1 Introduction

We study a class of procurement auctions under budget constraints. In a procurement auction, an auctioneer is interested in buying goods or services from a set of agents (also referred to as a reverse auction). The goal of the auctioneer is to select a subset of the agents so as to maximize her valuation function $v(\cdot)$, where $v(S)$ denotes the value derived when S is the selected subset of agents to get services from. The additional feature in the problem is that selecting an agent comes at a cost and there is a hard budget constraint that should not be violated. The algorithmic version then asks to maximize the generated value such that the total cost of the

selected agents should not exceed the budget, which results in natural "budgeted" versions of known optimization problems. Such problems are also motivated by recent application scenarios including crowdsourcing platforms, where agents can be viewed as workers providing tasks, e.g., [4], and influence maximization in social networks, where agents correspond to influential users [27].

Since the underlying optimization problem is NP-hard in most cases, our focus is on approximation algorithms. Furthermore, in the setting we consider here, the true cost of each agent is private information and we would like to design mechanisms that elicit truthful reporting by all agents. Hence, our ideal goal is to have truthful mechanisms that provide a good approximation to the optimal value for the auctioneer, and are *budget feasible*, i.e., the sum of the payments to the agents does not exceed the prespecified budget. Budget feasibility makes the problem more challenging, as it already rules out well known mechanisms such as VCG. We also stress that the question is nontrivial even if we allow exponential time algorithms, since computational power does not necessarily make the problem easier (see the discussion in [10]).

This framework of budget-feasible mechanism design was studied by Singer [26], when $v(\cdot)$ is an additive or a non-decreasing submodular function. Later on, follow up works also provided more results for richer classes of functions such as the class of XOS and subadditive functions (see the related work section). Most of these works however make the assumption that the valuation function is non-decreasing, i.e., $v(S) \leq v(T)$ for $S \subseteq T$, notable exceptions being the works of [10, 6]. Although monotonicity makes sense in several scenarios, one can think of examples where it is violated. E.g., if we consider influence maximization problems in social networks, adding more users to the selected set may some times bring negative influence [7] (some combinations of users may also not be compatible or well fitted together). Another prominent example of a non-monotone submodular function is the budgeted Max Cut problem, where $v(\cdot)$ is determined by the cuts of the graph. It is therefore natural to wonder if we can have truthful mechanisms with good performance for classes of non-monotone functions. Currently, the few results that are known for such problems have very large approximation ratios and often run in superpolynomial time. Even worse, in most cases, we do not even know of deterministic mechanisms (see Table 1).

Contribution: The main focus of our paper is on symmetric[1] submodular functions, where the value of a set S equals the value of its complement. This is an important class of non-monotone problems that contains the cut functions mentioned above, the mutual information function, and other examples that arise in optimization. We begin in Section 3 with a purely algorithmic result, obtaining a $\frac{2e}{e-1}$-approximation for symmetric submodular functions under a budget constraint. We believe this result is of independent interest, as it is the best known factor achieved by a deterministic algorithm (there exists already a randomized e-approximation). We then proceed to propose truthful, budget feasible mechanisms in Sections 4 and 5. Our results significantly improve the known approximation ratios for these problems, establishing at the same time polynomial running time for many cases where only exponential mechanisms were known. As an example, for the budgeted weighted cut problem we obtain the first deterministic polynomial time mechanism with a 27.25-approximation, and for unweighted cut functions we improve the approximation ratio for randomized mechanisms, from 564 down to 10. Analogous improvements are obtained also for arbitrary symmetric submodular functions. Finally, in Section 6 we briefly study the class of XOS functions, where we improve the current upper bound by a factor of 3.

[1] In some works on mechanism design, symmetric submodular functions have a different meaning and refer to the case where $v(S)$ depends only on $|S|$. Here we have adopted the terminology of earlier literature on submodular optimization, e.g., [13].

All our contributions in mechanism design are summarized in Table 1. We also stress that our mechanisms for general symmetric submodular functions use the value query model for oracle access to v, which is a much weaker requirement than the demand query model assumed in previous works.

Regarding the technical contribution of our work, the core idea of our approach is to exploit a combination of local search (either for finding exact or approximate local optima) with known mechanisms for non-decreasing submodular functions. The reason local search is convenient for symmetric submodular functions is that it produces two local optima, namely a set S and its complement, and we can then prove that the function $v(\cdot)$ is non-decreasing (or almost non-decreasing in a certain sense) within each local optimum. This allows us to utilize or adjust known mechanisms for non-decreasing submodular functions on the two subsets and then prove that one of the two solutions will attain a good approximation. The running time is still an issue under this approach, since finding an exact local optimum is not always guaranteed to terminate fast. However, by finding approximate local optima, and despite the fact that we lose some monotonicity properties this way, we are still able to obtain provable approximation guarantees.

	symmetric submod.		unweighted cut		weighted cut		XOS
	rand.	determ.	rand.	determ.	rand.	determ.	rand.
Previous work	768* [6]	–	564 [10]	5158 [10]	768* [6]	–	768* [6]
This paper	10*	10.90*	10	27.25	27.25	27.25	250*

Table 1: A summary of our results on mechanisms. The asterisk (*) indicates that the corresponding mechanism runs in superpolynomial time.

Related Work: The study of budget feasible mechanisms, as considered here, was initiated by Singer [26]. Later, Chen et al. [9] significantly improved Singer's results, obtaining a randomized 7.91-approximation mechanism and a deterministic 8.34-approximation mechanism for non-decreasing submodular functions. Several modifications of Chen et al.'s deterministic mechanism have been proposed that run in polynomial time for special cases [27, 17, 3]. For subadditive functions Dobzinski et al. [10] suggested a randomized $O(\log^2 n)$-approximation mechanism, and they gave the first constant factor mechanisms for non-monotone submodular objectives, specifically for *cut functions*. The factor for subadditive functions was later improved to $O\big(\frac{\log n}{\log\log n}\big)$ by Bei et al. [6], who also gave a randomized $O(1)$-approximation mechanism for XOS functions, albeit in exponential time, and further initiated the Bayesian analysis in this setting. Further improved $O(1)$-approximation mechanisms for XOS functions have also been obtained in [22]. Recently, Amanatidis et al. [3] suggested a 4-approximation mechanism for a subclass of XOS functions that captures matroid and matching constraints. There is also a line of related work under the *large market* assumption (where no participant can significantly affect the market outcome) which allows for polynomial time mechanisms with improved performance [28, 4, 14, 5, 18].

On maximization of submodular functions subject to knapsack or other type of constraints there is a vast literature, going back several decades, see, e.g., [24, 30]. More recently Lee et al. [21] provided the first constant factor randomized algorithm for submodular maximization under k matroid and k knapsack constraints, with factors $k+2+\frac{1}{k}$ and 5 respectively. The problem was also studied by Gupta et al. [16] who also proposed a randomized algorithm which achieves a

$(4+\alpha)$-approximation[2] in case of knapsack constraints where α is the approximation guarantee of the unconstrained submodular maximization. Later on Chekuri et al. [8] suggested a randomized 3.07-approximation algorithm improving the previously known results. Finally, Feldman et al. and Kulik et al. [12, 20] proposed their own randomized algorithms when there are knapsack constraints, achieving an e-approximation.

2 Notation and Preliminaries

We use $A = [n] = \{1, 2, ..., n\}$ to denote a set of n agents. Each agent i is associated with a private cost c_i, denoting the cost for participating in the solution. We consider a procurement auction setting, where the auctioneer is equipped with a valuation function $v : 2^A \to \mathbb{Q}_{\geq 0}$ and a budget $B > 0$. For $S \subseteq A$, $v(S)$ is the value derived by the auctioneer if the set S is selected (for singletons, we will often write $v(i)$ instead of $v(\{i\})$). Therefore, the algorithmic goal in all the problems we study is to select a set S that maximizes $v(S)$ subject to the constraint $\sum_{i \in S} c_i \leq B$. We assume oracle access to v via value queries, i.e., we assume the existence of a polynomial time value oracle that returns $v(S)$ when given as input a set S.

We mostly focus on a natural subclass of submodular valuation functions that includes *cut functions*, namely non-negative symmetric submodular functions. Throughout this work we make the natural assumption that $v(\emptyset) = 0$.

Definition 2.1. *A function v, defined on 2^A for some set A, is* submodular *if $v(S \cup \{i\}) - v(S) \geq v(T \cup \{i\}) - v(T)$ for any $S \subseteq T \subseteq A$, and $i \notin T$. Moreover, it is* non-decreasing *if $v(S) \leq v(T)$ for any $S \subseteq T \subseteq A$, while it is* symmetric *if $v(S) = v(A \setminus S)$ for any $S \subseteq A$.*

It is easy to see that v cannot be both symmetric and non-decreasing unless it is a constant function. In fact, if this is the case and $v(\emptyset) = 0$, then $v(S) = 0$, for all $S \subseteq A$. We also state an alternative definition of a submodular function, which will be useful later on.

Theorem 2.2 (Nemhauser et al. [24])**.** *A set function v is submodular if and only if for all $S, T \subseteq A$ we have $v(T) \leq v(S) + \sum_{i \in T \setminus S} (v(S \cup \{i\}) - v(S)) - \sum_{i \in S \setminus T} (v(S \cup T) - v(S \cup T \setminus \{i\}))$.*

We often need to argue about optimal solutions of sub-instances, from an instance we begin with. Given a cost vector $\mathbf{c}$, and a subset $X \subseteq A$, we denote by $\mathbf{c}_X$ the projection of $\mathbf{c}$ on X, and by $\mathbf{c}_{-X}$ the projection of $\mathbf{c}$ on $A \setminus X$. We also let $\text{OPT}(X, v, \mathbf{c}_X, B)$ be the value of an optimal solution to the restriction of this instance on X. Similarly, $\text{OPT}(X, v, \mathbf{c}_X, \infty)$ denotes the value of an optimal solution to the unconstrained version of the problem restricted on X. For the sake of readability, we often drop the valuation function and the cost vector, and write $\text{OPT}(X, B)$ or $\text{OPT}(X, \infty)$.

Finally, in Sections 3-5 we make one further assumption: we assume that there is at most one item whose cost exceeds the budget. As it is shown in Lemma A.2 and Corollary A.3 in Appendix A.2, this is without loss of generality.

Local Optima and Local Search. Given $v : 2^A \to \mathbb{Q}$, a set $S \subseteq A$ is called a $(1+\epsilon)$-*approximate local optimum* of v, if $(1+\epsilon)v(S) \geq v(S \setminus \{i\})$ and $(1+\epsilon)v(S) \geq v(S \cup \{i\})$ for every $i \in A$. When $\epsilon = 0$, S is called an *exact local optimum* of v. Note that if v is symmetric submodular, then S is a $(1+\epsilon)$-approximate local optimum if and only if $A \setminus S$ is a $(1+\epsilon)$-approximate local optimum.

[2] In the case of symmetric submodular functions the algorithm gives a deterministic 6-approximation.

Approximate local optima can produce good approximations in unconstrained maximization of non-monotone submodular functions [11], although here they are of interest for a quite different reason that becomes apparent in Lemmata 3.1 and 4.1. We can efficiently find approximate local optima using the local search algorithm APPROX-LOCAL-SEARCH of [11]. For completeness, the algorithm is included in Appendix A.3.

Lemma 2.3 (inferred from [11])**.** *Given a submodular function $v : 2^{[n]} \to \mathbb{Q}_{\geq 0}$,* APPROX-LOCAL-SEARCH$(A, v, \epsilon)$ *outputs a $(1 + \frac{\epsilon}{n^2})$-approximate local optimum using $O\left(\frac{1}{\epsilon} n^3 \log n\right)$ calls to the value oracle for v.*

2.1 Mechanism Design

In the strategic version that we consider here, every agent i only has his true cost c_i as private information. A mechanism $\mathcal{M} = (f, p)$ in our context consists of an outcome rule f and a payment rule p. Given a vector of cost declarations, $\mathbf{b} = (b_i)_{i \in A}$, where b_i denotes the cost reported by agent i, the outcome rule of the mechanism selects the set $f(\mathbf{b})$. At the same time, it computes payments $p(\mathbf{b}) = (p_i(\mathbf{b}))_{i \in N}$ where $p_i(\mathbf{b})$ denotes the payment issued to agent i. Hence, the final utility of agent i is $p_i(\mathbf{b}) - c_i$.

The properties we want to enforce in our mechanism design problem are the following.

Definition 2.4. *A mechanism $\mathcal{M} = (f, p)$ is*

1. truthful*, if reporting c_i is a dominant strategy for every agent i.*
2. individually rational*, if $p_i(\mathbf{b}) \geq 0$ for every $i \in A$, and $p_i(\mathbf{b}) \geq c_i$, for every $i \in f(\mathbf{b})$.*
3. budget feasible*, if $\sum_{i \in A} p_i(\mathbf{b}) \leq B$ for every $\mathbf{b}$.*

For randomized mechanisms, we use the notion of *universal truthfulness*, which means that the mechanism is a probability distribution over deterministic truthful mechanisms.

To design truthful mechanisms, we use the characterization by Myerson [23]. In particular, we say that an outcome rule f is *monotone*, if for every agent $i \in A$, and any vector of cost declarations $\mathbf{b}$, if $i \in f(\mathbf{b})$, then $i \in f(b'_i, \mathbf{b}_{-i})$ for $b'_i \leq b_i$. This simply means that if an agent i is selected by declaring cost b_i, then by declaring a lower cost he should still be selected.

Myerson's lemma (see Appendix A.4, Lemma A.1), implies that monotone algorithms admit truthful payment schemes (often referred to as *threshold payments*). This greatly simplifies the design of truthful mechanisms by focusing only on constructing monotone algorithms and not having to worry about the payment scheme. For all the mechanisms in this work, we assume that the underlying payment scheme is given by Myerson's lemma.

Mechanisms for Non-Decreasing Submodular Valuations. We will repeatedly make use of truthful budget feasible mechanisms for non-decreasing submodular valuation functions as subroutines in the mechanisms we design for non-monotone submodular functions. The best known such mechanisms are due to Chen et al. [9]. Here, we follow the improved analysis of [18] for the approximation ratio of the randomized mechanism RAND-MECH-SM of [9] (see Lemmata A.5 and D.2). With respect to this analysis, we also optimize the deterministic mechanism of [9] to obtain MECH-SM. Due to lack of space, these mechanisms are stated in Appendix A.4. The next theorem summarizes their properties and their performance.

Theorem 2.5 (inferred from [9], [18], and Lemma D.2)**.**

i. RAND-MECH-SM *runs in polynomial time, it is universally truthful, individually rational, budget-feasible, and has approximation ratio* 5.

ii. MECH-SM *is deterministic, truthful, individually rational, budget-feasible, and has approximation ratio* $3+\sqrt{6}$.

Similarly, in order to design deterministic mechanisms that run in polynomial time, we optimize a deterministic mechanism of [3] to obtain MECH-SM-FRAC. The difference with MECH-SM is that now we assume that a fractional relaxation of our problem can be solved optimally and that the fractional optimal solution is within a constant of the integral solution. Typically, this means that we are able to solve efficiently the LP relaxation of an ILP formulation of the problem, and to also bound the solution, often using pipage rounding.

Theorem 2.6 (inferred from [3] and Lemma D.2)**.** *Consider a relaxation of our problem for which we have an exact algorithm. Moreover, suppose that* $\text{OPT}_f(A, v, \mathbf{c}, B) \leq \rho \cdot \text{OPT}(A, v, \mathbf{c}, B)$ *for any instance, where* OPT_f *and* OPT *denote the value of an optimal solution to the relaxed and the original problem respectively. Then* MECH-SM-FRAC *is deterministic, truthful, individually rational, budget-feasible, and has approximation ratio* $\rho+2+\sqrt{\rho^2+4\rho+1}$. *Moreover, it runs in polynomial time given that the exact algorithm for the relaxed problem runs in polynomial time.*

Note that when the relaxed problem is the same as the original, then MECH-SM-FRAC becomes MECH-SM and the two approximation ratios coincide. A discussion on how different results combine into Theorems 2.5 and 2.6, is deferred to Appendix A.4.

Remark 2.7. In [6] the proposed mechanisms regard XOS and non-decreasing subadditive objectives, but it is stated that their results can be extended for general subadditive functions as well. This is achieved by defining $\hat{v}(S) = \max_{T \subseteq S} v(T)$. It is easily seen that $\hat{v}$ is non-decreasing, subadditive, and any solution that maximizes v is also an optimal solution for $\hat{v}$. Although this is true for subadditive functions, it does not hold for submodular functions. In particular, if v is submodular, then $\hat{v}$ is not necessarily submodular. Therefore the results of [9] cannot be extended to the setting that we study in this paper, even when time complexity is not an issue. An example of how $\hat{v}$ may fail to be submodular when v is a cut function is given in Appendix A.6, where we also discuss how it can be derived from [15] that in such cases $\hat{v}$ is actually XOS. This also implies that the best previously known approximation factor for the class of symmetric submodular functions was 768 (inherited by the use of $\hat{v}$).

3 The Core Idea: A Simple Algorithm for Symmetric Submodular Objectives

This section deals with the algorithmic version of the problem. The main result is a deterministic $\frac{2e}{e-1}$-approximation algorithm for symmetric submodular functions. Due to space constraints the proofs are deferred to Appendix B.

In this section the costs and the budget are assumed to be integral. Given a symmetric submodular function v, the goal is to find $S \subseteq A$ that maximizes $v(S)$ subject to the constraint $\sum_{i \in S} c_i \leq B$.

Since our function is not monotone, we cannot directly apply the result of [29], which gives an optimal simple greedy algorithm for non-decreasing submodular maximization subject to a knapsack constraint. Instead, our main idea is to combine appropriately the result of [29] with

the local search used for unconstrained symmetric submodular maximization [11]. On a high level, what happens is that local search produces an approximate solution S for the unconstrained problem, and while this is not our goal, v is "close to being non-decreasing" on both S and $A\setminus S$. This becomes precise in Lemma 3.1 below, but the point is that running a variant of the greedy algorithm of [29], on both S and $A\setminus S$ will now produce at least one good enough solution.

LS-GREEDY$(A, v, \mathbf{c}, B, \epsilon)$

$S =$ APPROX-LOCAL-SEARCH$(A, v, \epsilon/4)$
$T_1 =$ GREEDY-ENUM-SM$(S, v, \mathbf{c}, B)$
$T_2 =$ GREEDY-ENUM-SM$(A\setminus S, v, \mathbf{c}, B)$
Let T be the best solution among T_1 and T_2
return T

So the first component of our algorithm is the local search algorithm of [11] (Appendix A.3). By Lemma 2.3 and the fact that v is symmetric, both S and $A\setminus S$ are $(1+\frac{\epsilon}{4n^2})$-approximate local optima. We can now quantify the crucial observation that v is close to being non-decreasing within the approximate local optima S and $A\setminus S$. In fact we only need this property on the local optimum that contains the best feasible solution.

Lemma 3.1. *Let S be a $(1+\frac{\epsilon}{4n^2})$-approximate local optimum and consider*

$$X \in \operatorname*{arg\,max}_{Z\in\{S, A\setminus S\}} \text{OPT}(Z, v, \mathbf{c}_Z, B)\,.$$

Then, for every $T \subsetneq X$ and every $i \in X\setminus T$, we have $v(T\cup\{i\}) - v(T) > -\frac{\epsilon}{n}\,\text{OPT}(X, B)$.

The second component of LS-GREEDY is an appropriate modification of the greedy algorithm of [29] for non-monotone submodular functions. It first enumerates all solutions of size at most 3. Then, starting from each 3-set, it builds a greedy solution, and it outputs the best among these $\Theta(n^3)$ solutions. Here this idea is adjusted for non-monotone functions.

GREEDY-ENUM-SM$(A, v, \mathbf{c}, B)$

Let S_1 be the best feasible solution of cardinality at most 3 (by enumerating all such solutions)
$S_2 = \emptyset$
for *every $U \subseteq A$ with $|U| = 3$* **do**
    $T = U,\ t = 1,\ A^0 = A\setminus U$
    while $A^{t-1} \neq \emptyset$ **do**
        Find $\theta_t = \max_{i\in A^{t-1}} \dfrac{v(T\cup\{i\}) - v(T)}{c_i}$, and let i_t be an element of A^{t-1} that attains θ_t
        if $\theta_t \geq 0$ *and* $\sum_{i\in T\cup\{i_t\}} c_i \leq B$ **then**
            $T = T\cup\{i_t\}$
        $A^t = A^{t-1}\setminus\{i_t\}$
        $t = t+1$
    if $v(T) > v(S_2)$ **then**
        $S_2 = T$
Let S be the best solution among S_1 and S_2
return S

As explained in Fact B.1 in Appendix B, at least one of S and $A\setminus S$ contains a feasible solution of value at least $0.5\,\text{OPT}(A, B)$. Lemma 3.1 guarantees that in this set, v is very close to a non-

decreasing submodular function. This is sufficient for GREEDY-ENUM-SM to perform almost as well as if v was non-decreasing.

Theorem 3.2. *For any $\epsilon > 0$, algorithm* LS-GREEDY *achieves a $\left(\frac{2e}{e-1}+\epsilon\right)$-approximation.*

Theorem 3.2 suggests that a straightforward composition of two well known greedy algorithms achieves a good approximation for any symmetric submodular objective. We believe this is of independent interest and could be useful for other problems involving submodular optimization. From a mechanism design perspective, however, algorithm LS-GREEDY fails to be monotone and thus it cannot be used directly in the subsequent sections. In the next two sections, we remedy the problem by removing the enumeration part of the algorithm.

4 Mechanisms for Symmetric Submodular Objectives: A First Take

Utilizing the algorithmic approach of Section 3 to get truthful mechanisms is not straightforward. The reason is not only that LS-GREEDY is not monotone. The greedy algorithm without the enumeration part *is* monotone even for general objectives, but, to further complicate things, it is not guaranteed to be budget-feasible or have a good performance anymore. Instead of computing approximate local optima like in Section 3, in this section we bypass most issues by computing exact local optima. The highlights of this simplified approach are polynomial mechanisms for unweighted cut functions with greatly improved guarantees.

The price we have to pay however, is that in general, finding exact local optima is not guaranteed to run in polynomial time [25]. Still, these mechanisms deepen our understanding of the problem. As mentioned in the Introduction, the problem seems to remain hard even when the running time is not an issue, and many existing mechanisms for various classes of functions are not polynomial. In particular, there are no better known mechanisms—even running in exponential time—for symmetric submodular objectives. We are going to deal further with the issue of running time in Section 5.

Below we give a randomized mechanism that reduces the known factor of 768 down to 10, as well as the first deterministic $O(1)$-approximation mechanism for symmetric submodular objectives. In both mechanisms, local search produces a local maximum S for the unbudgeted problem and then the budgeted problem is solved optimally on both S and $A\setminus S$. As shown in Lemma 4.1, v is non-decreasing on both S and $A\setminus S$. Thus, running the mechanisms of [9] on $T \in \arg\max_{X\in\{S,A\setminus S\}} \text{OPT}(X,B)$ directly implies a good solution. Since they are very similar, we state the two mechanisms together.

RAND-MECH-SYMSM$(A, v, \mathbf{c}, B)$ (*resp.* DET-MECH-SYMSM$(A, v, \mathbf{c}, B)$)

$S =$ APPROX-LOCAL-SEARCH$(A, v, 0)$ //find an exact local optimum
if OPT$(S, v, \mathbf{c}_S, B) \geq$ OPT$(A\setminus S, v, \mathbf{c}_{A\setminus S}, B)$ **then**
  return RAND-MECH-SM$(S, v, \mathbf{c}_S, B)$ (*resp.* MECH-SM$(S, v, \mathbf{c}_S, B)$)
else
  return RAND-MECH-SM$(A\setminus S, v, \mathbf{c}_{A\setminus S}, B)$ (*resp.* MECH-SM$(A\setminus S, v, \mathbf{c}_{A\setminus S}, B)$)

The next simple lemma is crucial for the performance of both mechanisms for arbitrary submodular functions, and it shows how local search helps us exploit known results for non-decreasing submodular functions.

Lemma 4.1. *Let A be a set and v be an arbitrary submodular function defined on 2^A. If S is a local maximum of v, then v is submodular and non-decreasing when restricted on 2^S.*

Since v is symmetric, if S is a local optimum, so does $A \setminus S$. Lemma 4.1 suggests that we can use the mechanism RAND-MECH-SM (resp. MECH-SM) on S and $A \setminus S$, to get the following implications. The proof of the theorem is in Appendix C.

Theorem 4.2. *i. The mechanism* RAND-MECH-SYMSM *is universally truthful, individually rational, budget-feasible, and has approximation ratio* 10.
ii. The mechanism DET-MECH-SYMSM *is deterministic, truthful, individually rational, budget-feasible, and has approximation ratio* $6 + 2\sqrt{6}$.

Lower bounds on the approximability have been obtained by Chen et al. [9] for additive valuations. Since additive functions are not symmetric, these lower bounds do not directly apply here. However, it is not hard to construct symmetric submodular functions that give the exact same bounds. The following is established in Appendix C.

Fact 4.3. *Independent of complexity assumptions, there is no deterministic (resp. randomized) truthful, budget feasible mechanism that can achieve an approximation ratio better than* $1 + \sqrt{2}$ *(resp.* 2*).*

Clearly, both mechanisms require superpolynomial time in general, unless $P = NP$. In both cases, instead of $\text{OPT}(\cdot, B)$ we could use the optimal solution of a fractional relaxation of the problem, at the expense of somewhat worse guarantees. This does not completely resolve the problem, although this way local search becomes the sole bottleneck. For certain objectives, however, we can achieve similar guarantees in polynomial time. Unweighted cut functions are the most prominent such example, and it is the focus of the next subsection.

4.1 Polynomial Time Mechanisms for Unweighted Cut Functions

We begin with the definition of the problem when v is a cut function:

Budgeted Max Weighted Cut. Given a complete graph G with vertex set $V(G) = [n]$, nonnegative weights w_{ij} and costs c_{ij} on the edges, and a positive budget B, find $X \subseteq [n]$ so that $v(X) = \sum_{i \in X} \sum_{j \in [n] \setminus X} w_{ij}$ is maximized subject to $\sum_{j \in X} c_j \leq B$.

For convenience, we assume the problem is defined on a complete graph as we can take $w_{ij} = 0$ to incorporate any graphs. In this subsection, we focus on the unweighted version, where all weights are equal to either 0 or 1. We call this special case *Budgeted Max Cut.* The weighted version is considered in Subsection 5.1.

The fact that local search takes polynomial time to find an exact local optimum for the unweighted version [19] does not suffice to make RAND-MECH-SYMSM a polynomial time mechanism, since one still needs to compute $\text{OPT}(S, B)$ and $\text{OPT}(A \setminus S, B)$. However, a small modification so that RAND-MECH-SM(S, B) and RAND-MECH-SM$(A \setminus S, B)$ are returned with probability 1/2 each (see Appendix D) yields a randomized 10-approximate polynomial time mechanism. This implies the following theorem.

Theorem 4.4. *There is a universally truthful, individually rational, budget-feasible mechanism for Budgeted Max Cut that has approximation ratio* 10 *and runs in polynomial time.*

To obtain a deterministic mechanism, we will use an LP-based approach. In particular, we will need first to compare the value of an optimal solution of a fractional relaxation to the value of an optimal solution of the original problem.

Ageev and Sviridenko [1] studied a different Max Cut variant, but we follow a similar approach to obtain the desired bound for Budgeted Max Weighted Cut. We begin with a linear program formulation of the problem. Our analysis is carried out for the weighted version of the problem since we are going to reuse some results in Subsection 5.1 which deals with weighted cut functions. In fact, we want to argue about a variant of the problem where we may only be allowed to chose the solution from a specified subset of A. That is, we formulate the sub-instance $I = (X, v, \mathbf{c}_X, B)$ of $(A, v, \mathbf{c}, B)$

We associate a binary variable x_i for each vertex i, and the partition of $A = [n]$ according to the value of the x_is defines the cut. There is also a binary variable z_{ij} for each edge $\{i, j\}$ which is the indicator variable of whether $\{i, j\}$ is in the cut.

$$\text{maximize:} \sum_{i\in[n]} \sum_{j\in[n]\setminus[i]} w_{ij} z_{ij} \tag{1}$$

$$\text{subject to:} \quad z_{ij} \le x_i + x_j\,, \qquad \forall i \in [n], \forall j \in [n]\setminus[i] \tag{2}$$

$$z_{ij} \le 2 - x_i - x_j\,, \qquad \forall i \in [n], \forall j \in [n]\setminus[i] \tag{3}$$

$$\sum_{i\in[n]} c_i x_i \le B \tag{4}$$

$$x_i = 0\,, \qquad \forall i \in [n]\setminus X \tag{5}$$

$$0 \le x_i, z_{ij} \le 1\,, \qquad \forall i \in [n], \forall j \in [n]\setminus[i] \tag{6}$$

$$x_i \in \{0, 1\}\,, \qquad \forall i \in [n] \tag{7}$$

It is not hard to see that (1)-(7) is a natural ILP formulation for Budgeted Max Weighted Cut and (1)-(6) is its linear relaxation. Let $\text{OPT}(I)$ and $\text{OPT}_f(I)$ denote the optimal solutions to (1)-(7) and (1)-(6) respectively for instance I. To show how these two are related we will use the technique of pipage rounding [1, 2]. Although we do not provide a general description of the technique, the proof of the next theorem is self-contained (see Appendix C).

Theorem 4.5. *Given the fractional relaxation* (1)-(6) *for Budgeted Max Weighted Cut, we have that* $\text{OPT}_f(I) \le (2 + 2\beta_I) \cdot \text{OPT}(I)$*, for any instance* I*, where* β_I *is such that* $\max_{i\in A} v(i) \le \beta_I \cdot \text{OPT}(I)$.

Note that $\beta_I \le 1$ for every instance I, hence, the above theorem implies a worst-case upper bound of 4. Now, we may modify DET-MECH-SYMSM to use OPT_f instead of OPT, and MECH-SM-FRAC instead of MECH-SM. This results in the following deterministic mechanism that runs in polynomial time.

DET-MECH-UCUT$(A, v, \mathbf{c}, B)$

```
1 Set A' = {i | c_i ≤ B} and i* ∈ argmax_{i∈A'} v(i)
2 if 26.25 · v(i*) ≥ OPT_f(A\{i*}, v, c_{-i*}, B) then
3     return i*
4 else
5     S = APPROX-LOCAL-SEARCH(A, v, 0)
6     if OPT_f(S, v, c_S, B) ≥ OPT_f(A\S, v, c_{A\S}, B) then
7         return MECH-SM-FRAC(S, v, c_S, B)
8     else
9         return MECH-SM-FRAC(A\S, v, c_{A\S}, B)
```

Theorem 4.6. DET-MECH-UCUT *is a deterministic, truthful, individually rational, budget-feasible mechanism for Budgeted Max Cut that has approximation ratio* 27.25 *and runs in polynomial time.*

5 Mechanisms for Symmetric Submodular Objectives Revisited

Suppose that for a symmetric submodular function v, an optimal fractional solution can be found efficiently and that $\text{OPT}_f(A,B) \leq \rho \cdot \text{OPT}(A,B)$ for any instance, where OPT_f and OPT denote the value of an optimal solution to the relaxed and the original problem respectively.

A natural question is whether the approach taken for unweighted cut functions can be fruitful for other symmetric submodular objectives. In the mechanisms of Subsection 4.1, however, the complexity of local search can be a bottleneck even for objectives where an optimal fractional solution can be found fast and it is not far from the optimal integral solution. So, we now return to the idea of Section 3, where local search runs in polynomial time and produces an approximate local maximum; unfortunately, the nice property of monotonicity in each side of the partition (Lemma 4.1) does not hold any longer.

This means that the approximation guarantees of such mechanisms do not follow in any direct way from existing work. Moreover, budget-feasibility turns out to be an even more delicate issue since it crucially depends on the (approximate) monotonicity of the valuation function. Specifically, when a set X only contains a very poor solution to the original problem, every existing proof of budget feasibility for the restriction of v on X completely breaks down. Since we cannot expect that an approximate local maximum S and its complement $A \setminus S$ both contain a "good enough" solution to the original problem, we need to make sure that GREEDY-SM never runs on the wrong set.

The mechanism DET-MECH-UCUT for the unweighted cut problem seems to take care of this and we are going to build on it, in order to propose mechanisms for arbitrary symmetric submodular functions. To do so we replace the constant 26.25 that appears in DET-MECH-UCUT by $\alpha = (1+\rho)\left(2+\rho+\sqrt{\rho^2+4\rho+1}\right)-1$ and we find an approximate local maximum instead of an exact local maximum. Most importantly, in order to achieve budget-feasibility we introduce a modification of MECH-SM-FRAC (Appendix D.2) that runs GREEDY-SM with a slightly reduced budget. The ϵ' is determined by the analysis of the mechanism and only depends on the constants ρ and ϵ.

DET-MECH-SYMSM-FRAC$(A, v, \mathbf{c}, B)$

Set $A' = \{i \mid c_i \leq B\}$ and $i^* \in \arg\max_{i \in A'} v(i)$
if $\alpha \cdot v(i^*) \geq \text{OPT}_f(A \setminus \{i^*\}, v, \mathbf{c}_{-i^*}, B)$ **then**
    return i^*
else
    $S =$ APPROX-LOCAL-SEARCH(A, v, ϵ')
    if $\text{OPT}_f(S, v, \mathbf{c}_S, B) \geq \text{OPT}_f(A \setminus S, v, \mathbf{c}_{A \setminus S}, B)$ **then**
        return MECH-SM-FRAC-VAR$(S, v, \mathbf{c}_S, B, (1-(\alpha_1+2)\epsilon'))$
    else
        return MECH-SM-FRAC-VAR$(A \setminus S, v, \mathbf{c}_{A \setminus S}, B, (1-(\alpha_1+2)\epsilon'))$

Theorem 5.1 below shows that for any objective for which we can establish a constant upper bound ρ on the ratio of the fractional and the integral optimal solutions, we have constant factor

approximation mechanisms that run in polynomial-time .

Theorem 5.1. *For any $\epsilon > 0$,* DET-MECH-SYMSM-FRAC *is a deterministic, truthful, individually rational, budget-feasible mechanism for symmetric submodular valuations, that has approximation ratio $\alpha + 1 + \epsilon$ and runs in polynomial time.*

There are several steps towards proving the theorem (Appendix D), since we need the good properties of GREEDY-SM to still hold even for objectives that are not exactly non-decreasing.

5.1 Weighted Cut Functions

Let us return now to the Max Cut problem, and consider the weighted version. An immediate implication of Theorem 5.1 is that we get a deterministic polynomial-time mechanism for Budgeted Max Weighted Cut with approximation ratio 58.72. These is just the result of substituting $\rho = 4$—as suggested by Theorem 4.5—in the formula for α.

Corollary 5.2. *There is a deterministic, truthful, individually rational, budget-feasible mechanism that has approximation ratio 58.72 and runs in polynomial time.*

However, Theorem 4.5 says something stronger: given that $\max_{i \in A} v(i)$ is small compared to $\text{OPT}(A, B)$, ρ is strictly smaller than 4. Note that the first step in DET-MECH-SYMSM-FRAC is to compare $\max_{i \in A} v(i)$ to $\text{OPT}_f(A \setminus \{i^*\}, B)$. This implies an upper bound on $\max_{i \in A} v(i)$ in the following steps and we can use it to further fine-tune our mechanism. In particular, by setting $\alpha_1 = 26.245$ instead of $(1+4)\left(2+4+\sqrt{16+16+1}\right) - 1 = 57.72$ in DET-MECH-SYMSM-FRAC, we prove in Appendix D.2 the following improved result that matches the approximation guarantee for unweighted cut functions.

Theorem 5.3. *There is a deterministic, truthful, individually rational, budget-feasible mechanism for Budgeted Max Weighted Cut that has approximation ratio 27.25, and runs in polynomial time.*

6 An Improved Upper Bound for XOS Objectives

In [6] a randomized, universally truthful and budget-feasible 768-approximation mechanism was introduced for XOS functions. For several of the results in Sections 4 and 5 the best previously known upper bound follow from this result. In this section we slightly modify their mechanism to improve its performance.[3] We begin with the definition of XOS functions.

Definition 6.1. *A valuation function, defined on 2^A for some set A, is* XOS *or* fractionally subadditive*, if there exist non-negative additive functions $\alpha_1, \alpha_2, ..., \alpha_r$, for some finite r, such that $v(S) = \max\{\alpha_1(S), \alpha_2(S), ..., \alpha_r(S)\}$.*

In Appendix E we provide mechanism MAIN-XOS which modifies the parameters of the mechanism of Bei et al. By following a similar but more careful analysis, we improve the approximation ratio by a factor of 3 while also retaining its properties. Notice that like the mechanism of Bei et al., MAIN-XOS is randomized and has superpolynomial running time. In particular it requires a demand oracle.

Theorem 6.2. MAIN-XOS *is universally truthful, individually rational, budget-feasible, and has approximation ratio* 250.

[3] In a recent unpublished manuscript Leonardi et al. [22] also suggest a fine-tuning of this mechanism that yields a factor 436.

7 Conclusions

In this work, we have made progress on the design of budget-feasible mechanisms for symmetric submodular objectives. The highlights of our results are polynomial time algorithms for the Budgeted Max Cut problem (weighted and unweighted) with significant improvements over previously known approximation factors. Although for general symmetric submodular functions we have exponential running times, the results imply polynomial time algorithms for any objective, where we can bound the optimal fractional solution with respect to the integral one. We also expect that the idea of utilizing local search in order to identify monotone regions of a general submodular function may have further applications on mechanism design. Finally, apart from the mechanism design problem, the result of Section 3 is an interesting consequence for submodular optimization.

References

[1] Alexander A. Ageev and Maxim Sviridenko. Approximation algorithms for maximum coverage and max cut with given sizes of parts. In *Integer Programming and Combinatorial Optimization, 7th International IPCO Conference, Graz, Austria, June 9-11, 1999, Proceedings*, pages 17–30, 1999. doi:10.1007/3-540-48777-8_2.

[2] Alexander A. Ageev and Maxim Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3):307–328, 2004. doi:10.1023/B:JOCO.0000038913.96607.c2.

[3] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. Coverage, matching, and beyond: New results on budgeted mechanism design. In *Proceedings of the 12th International Conference on Web and Internet Economics, WINE 2016*, pages 414–428, 2016.

[4] Nima Anari, Gagan Goel, and Afshin Nikzad. Mechanism design for crowdsourcing: An optimal 1-1/e competitive budget-feasible mechanism for large markets. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 266–275, 2014. doi:10.1109/FOCS.2014.36.

[5] Eric Balkanski and Jason D. Hartline. Bayesian budget feasibility with posted pricing. In *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11 - 15, 2016*, pages 189–203. ACM, 2016.

[6] Xiaohui Bei, Ning Chen, Nick Gravin, and Pinyan Lu. Budget feasible mechanism design: from prior-free to bayesian. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 449–458, 2012. doi:10.1145/2213977.2214020.

[7] Allan Borodin, Yuval Filmus, and Joel Oren. Threshold models for competitive influence in social networks. In *Proceedings of the 6th International Workshop on Internet and Network Economics, WINE 2010*, pages 539–550, 2010.

[8] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM J. Comput.*, 43(6):1831–1879, 2014.

[9] Ning Chen, Nick Gravin, and Pinyan Lu. On the approximability of budget feasible mechanisms. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*

2011, San Francisco, California, USA, January 23-25, 2011, pages 685–699, 2011. doi:10.1137/1.9781611973082.54.

[10] Shahar Dobzinski, Christos H. Papadimitriou, and Yaron Singer. Mechanisms for complement-free procurement. In *Proceedings 12th ACM Conference on Electronic Commerce (EC-2011), San Jose, CA, USA, June 5-9, 2011*, pages 273–282, 2011. doi:10.1145/1993574.1993615.

[11] Uriel Feige, Vahab S. Mirrokni, and Jan Vondrák. Maximizing non-monotone submodular functions. *SIAM J. Comput.*, 40(4):1133–1153, 2011.

[12] Moran Feldman, Joseph Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 570–579. IEEE Computer Society, 2011.

[13] S. Fujishige. Canonical decompositions of symmetric submodular systems. *Discrete Applied Mathematics*, 5:175–190, 1983.

[14] Gagan Goel, Afshin Nikzad, and Adish Singla. Mechanism design for crowdsourcing markets with heterogeneous tasks. In *Proceedings of the Seconf AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2014, November 2-4, 2014, Pittsburgh, Pennsylvania, USA*, 2014.

[15] Anupam Gupta, Viswanath Nagarajan, and Sahil Singla. Adaptivity gaps for stochastic probing: Submodular and XOS functions. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1688–1702. SIAM, 2017.

[16] Anupam Gupta, Aaron Roth, Grant Schoenebeck, and Kunal Talwar. Constrained non-monotone submodular maximization: Offline and secretary algorithms. In *Internet and Network Economics - 6th International Workshop, WINE 2010, Stanford, CA, USA, December 13-17, 2010. Proceedings*, volume 6484 of *Lecture Notes in Computer Science*, pages 246–257. Springer, 2010.

[17] Thibaut Horel, Stratis Ioannidis, and S. Muthukrishnan. Budget feasible mechanisms for experimental design. In *LATIN 2014: Theoretical Informatics - 11th Latin American Symposium, Montevideo, Uruguay, March 31 - April 4, 2014. Proceedings*, pages 719–730, 2014. doi:10.1007/978-3-642-54423-1_62.

[18] Pooya Jalaly and Eva Tardos. Simple and Efficient Budget Feasible Mechanisms for Monotone Submodular Valuations, arxiv:1703:10681, 2017.

[19] Jon Kleinberg and Eva Tardos. *Algorithm Design*. Addison Wesley, 2006.

[20] Ariel Kulik, Hadas Shachnai, and Tami Tamir. Approximations for monotone and nonmonotone submodular maximization with knapsack constraints. *Math. Oper. Res.*, 38(4):729–739, 2013.

[21] Jon Lee, Vahab S. Mirrokni, Viswanath Nagarajan, and Maxim Sviridenko. Maximizing nonmonotone submodular functions under matroid or knapsack constraints. *SIAM J. Discrete Math.*, 23(4):2053–2078, 2010.

[22] Stefano Leonardi, Gianpiero Monaco, Piotr Sankowski, and Qiang Zhang. Budget Feasible Mechanisms on Matroids, arXiv:1612:03150, 2016.

[23] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1), 1981.

[24] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions - I. *Math. Program.*, 14(1):265–294, 1978.

[25] Alejandro A. Schäffer and Mihalis Yannakakis. Simple local search problems that are hard to solve. *SIAM J. Comput.*, 20(1):56–87, 1991.

[26] Yaron Singer. Budget feasible mechanisms. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 765–774, 2010. doi:10.1109/FOCS.2010.78.

[27] Yaron Singer. How to win friends and influence people, truthfully: influence maximization mechanisms for social networks. In *Proceedings of the Fifth International Conference on Web Search and Web Data Mining, WSDM 2012, Seattle, WA, USA, February 8-12, 2012*, pages 733–742, 2012. doi:10.1145/2124295.2124381.

[28] Adish Singla and Andreas Krause. Incentives for privacy tradeoff in community sensing. In *Proceedings of the First AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2013, November 7-9, 2013, Palm Springs, CA, USA*. AAAI, 2013.

[29] Maxim Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Oper. Res. Lett.*, 32(1):41–43, 2004.

[30] Laurence A. Wolsey. Maximising real-valued submodular functions: Primal and dual heuristics for location problems. *Math. Oper. Res.*, 7(3):410–425, 1982.

A Missing Material from Section 2

A.1 Myerson's Lemma

Here we state Myerson's lemma [23], which we use in all the mechanisms we propose. Recall that an outcome rule f is monotone, if for every agent $i \in A$, and any vector of cost declarations $\mathbf{b}$, if $i \in f(\mathbf{b})$, then $i \in f(b_i', \mathbf{b}_{-i})$ for $b_i' \leq b_i$.

Lemma A.1. *Given a monotone algorithm f, there is a unique payment scheme p such that (f, p) is a truthful and individually rational mechanism, given by*

$$p_i(\mathbf{b}) = \begin{cases} \sup_{b_i \in [c_i, \infty)} \{b_i : i \in f(b_i, \mathbf{b}_{-i})\}, & \text{if } i \in f(\mathbf{b}) \\ 0, & \text{otherwise} \end{cases}$$

Given Myerson's lemma, to obtain a truthful mechanism all we need to ensure is that the outcome rule is monotone.

A.2 Instances with Costs Exceeding the Budget

Consider an instance with a symmetric submodular function, where there exist agents with $c_i > B$. It may seem at first sight that we could just discard such agents, since too expensive agents are not included in any feasible solution anyway. However, the presence of such agents can create infeasible solutions of very high value and make an analog of Lemma 3.1 impossible to prove. Simply discarding them could also destroy the symmetry of the function (e.g., if we had a cut function defined on a graph, we could not just remove a node).

Let $\mathscr{I}$ denote the set of all instances of the problem with symmetric submodular functions, and let $\mathscr{J}$ denote the set of all instances where at most one agent has cost more than B. Given $X \subseteq A$, we let $\mathbf{c}(X) = \sum_{i \in X} c_i$. The next lemma, together with its corollary, show that, when dealing

with symmetric submodular functions, we may only consider instances in $\mathscr{J}$ without any loss of generality.

Lemma A.2. *Given an instance $I = (A, v, \mathbf{c}, B) \in \mathscr{I}$, we can efficiently construct an instance $J = (A', v', \mathbf{c}', B) \in \mathscr{J}$ such that*

- *Every feasible solution of I is a feasible solution of J and vice versa.*
- *If X is a feasible solution of I, then $v(X) = v'(X)$ and $\mathbf{c}(X) = \mathbf{c}'(X)$. In particular,* $\textsc{opt}(J) = \textsc{opt}(I)$.

Proof. Let $E = \{i \in A \mid c_i > B\}$ be the set of expensive agents and define $A' = (A \setminus E) \cup \{i_E\}$, where i_E is a new agent replacing the whole set E. For $i \in A \setminus E$ we define $c'(i) = c(i)$, while $c'(i_E) = B + 1$. Finally, v' is defined as follows

$$v'(T) = \begin{cases} v(T), & \text{if } T \subseteq A \setminus E \\ v((T \setminus \{i_E\}) \cup E), & \text{otherwise} \end{cases}$$

Now suppose X is a budget-feasible solution of I. Then $\mathbf{c}(X) \leq B$ and thus $X \subseteq A \setminus E$. But then, by the definition of $\mathbf{c}'$, $\mathbf{c}'(X) = \mathbf{c}(X) \leq B$ as well, and therefore X is also a budget-feasible solution of J. Moreover, $v'(X) = v(X)$ by the definition of v'. We conclude that $\textsc{opt}(I) \leq \textsc{opt}(J)$.

The proof that every feasible solution of J is a feasible solution of I is almost identical. This implies $\textsc{opt}(J) \leq \textsc{opt}(I)$, and therefore $\textsc{opt}(J) = \textsc{opt}(I)$. □

Now, it is not hard to see that we can turn any algorithmic result on $\mathscr{J}$ to the same algorithmic result on $\mathscr{I}$. However, we need a somewhat stronger statement to take care of issues like truthfulness and budget-feasibility. This is summarized in the following corollary.

Corollary A.3. *Given a (polynomial time) algorithm* $\textsc{alg}'$ *that achieves a ρ-approximation on instances in $\mathscr{J}$, we can efficiently construct a (polynomial time) ρ-approximation algorithm* ALG *that works for all instances in $\mathscr{I}$. Moreover, if* $\textsc{alg}'$ *is monotone and budget-feasible on instances in $\mathscr{J}$, assuming Myerson's threshold payments, then* ALG *is monotone and budget-feasible on instances in $\mathscr{I}$.*

Proof. The description of ALG is quite straightforward. Given an instance $I = (A, v, \mathbf{c}, B) \in \mathscr{I}$, ALG first constructs instance $J = (A', v', \mathbf{c}', B) \in \mathscr{J}$, as described in the proof of Lemma A.2. Then ALG runs $\textsc{alg}'$ with input J and returns its output. Clearly, if $\textsc{alg}'$ runs in polynomial time, so does ALG.

If $X = \textsc{alg}'(J) = \textsc{alg}(I)$, then X is feasible with respect to J and $\textsc{opt}(J) \leq \rho \cdot v'(X)$. By Lemma A.2 we get that X is feasible with respect to I and $\textsc{opt}(I) = \textsc{opt}(J) \leq \rho \cdot v'(X) = \rho \cdot v(X)$. This establishes the approximation ratio of ALG.

Next, assume that $\textsc{alg}'$ is monotone and—assuming Myerson's threshold payments—budget-feasible on instances in $\mathscr{J}$. Suppose that agent $j \in \textsc{alg}(I)$ reduces his cost from c_j to $b_j < c_j$. This results in a new instance $I_* = (A, v, (b_j, \mathbf{c}_{-j}), B) \in \mathscr{I}$. Since it must be the case where $c_j \leq B$, the corresponding instance of $\mathscr{J}$ is $J_* = (A', v', (b_j, \mathbf{c}'_{-j}), B)$. Due to the monotonicity of $\textsc{alg}'$ we have

$$j \in \textsc{alg}(I) = \textsc{alg}'(J) \Rightarrow j \in \textsc{alg}'(J_*) = \textsc{alg}(I_*),$$

and therefore ALG is monotone as well.

The budget-feasibility of ALG follows from the budget-feasibility of $\textsc{alg}'$ by observing that $i \in \textsc{alg}(A, v, (b_j, \mathbf{c}_{-j}), B)$ if and only if $i \in \textsc{alg}'(A', v', (b_j, \mathbf{c}_{-j}), B)$ for all $i \in A'$. □

A.3 A Local Search Algorithm

We present below the algorithm of [11] for computing approximate local optima.

APPROX-LOCAL-SEARCH(A, v, ϵ) [11]

$S = \{i^*\}$, where $i^* \in \arg\max_{i \in A} v(i)$
while *there exists some a such that* $\max\{v(S \cup \{a\}), v(S \setminus \{a\})\} > (1 + \epsilon/n^2)\, v(S)$ **do**
    if $v(S \cup \{a\}) > (1 + \epsilon/n^2)\, v(S)$ **then**
        $S = S \cup \{a\}$
    else
        $S = S \setminus \{a\}$
return S

A.4 Mechanisms for Non-Decreasing Submodular Valuation Functions

For non-decreasing submodular valuations, Chen et al. [9] give a randomized, budget feasible, individually rational, truthful mechanism that achieves a constant approximation.

RAND-MECH-SM$(A, v, \mathbf{c}, B)$ [9]

Set $A' = \{i \mid c_i \leq B\}$ and $i^* \in \arg\max_{i \in A'} v(i)$
with probability $\frac{2}{5}$ **return** i^*
with probability $\frac{3}{5}$ **return** GREEDY-SM$(A, v, \mathbf{c}, B/2)$

The mechanism GREEDY-SM is a greedy algorithm that picks agents according to their ratio of marginal value over cost, given that this cost is not too large. For the sake of presentation, we assume the agents are sorted in descending order with respect to this ratio. The marginal value of each agent is calculated with respect to the previous agents in the ordering, i.e. $1 = \arg\max_{j \in A} \frac{v(j)}{c_j}$ and $i = \arg\max_{j \in A \setminus [i-1]} \frac{v([j]) - v([j-1])}{c_j}$ for $i \geq 2$.

GREEDY-SM$(A, v, \mathbf{c}, B/2)$ [9]

Let $k = 1$ and $S = \emptyset$
while $k \leq |A|$ **and** $v(S \cup \{k\}) > v(S)$ **and** $c_k \leq \frac{B}{2} \cdot \frac{v(S \cup \{k\}) - v(S)}{v(S \cup \{k\})}$ **do**
    $S = S \cup \{k\}$
    $k = k + 1$
return S

Lemma A.4 (inferred from [9] and [18])**.** GREEDY-SM *is monotone. Assuming a non-decreasing submodular function v and the payments of Myerson's lemma,* GREEDY-SM$(A, v, \mathbf{c}, B/2)$ *is budget-feasible and outputs a set S such that $v(S) \geq \frac{1}{3} \cdot \text{OPT}(A, v, \mathbf{c}, B) - \frac{2}{3} \cdot v(i^*)$.*

A derandomized version of RAND-MECH-SM is also provided. It has all the desired properties, while suffering a small loss on the approximation factor. Here, following the improved analysis of [18] for the ratio of RAND-MECH-SM, we fine-tune this derandomized mechanism to obtain MECH-SM that has a better approximation guarantee (see Theorem 2.5 and the following subsection).

MECH-SM$(A, v, \mathbf{c}, B)$

Set $A' = \{i \mid c_i \leq B\}$ and $i^* \in \arg\max_{i \in A'} v(i)$
if $(2+\sqrt{6}) \cdot v(i^*) \geq \text{OPT}(A \setminus \{i^*\}, v, \mathbf{c}_{-i^*}, B)$ **then**
    return i^*
else
    return GREEDY-SM$(A, v, \mathbf{c}, B/2)$

In both cases, either an agent i^* of maximum value or GREEDY-SM$(A, v, \mathbf{c}, B/2)$ is returned. In RAND-MECH-SM this is done according to a probability distribution, while in MECH-SM the decision depends on the comparison of i^* with an optimal solution at the instance $A \setminus \{i^*\}$ with budget B.

In [3], a variant of MECH-SM was used for the case where a relaxation of the problem can be solved optimally and its solution is within a constant ρ of the optimal solution of the original problem. There, as well as in [27], the LP relaxation of an ILP formulation of the problem was used for coverage objectives. Here we optimize this mechanism to get MECH-SM-FRAC below (see Theorem 2.6 and the following subsection). By OPT_f we denote the optimal (fractional) solution of the relaxed budgeted problem. Notice that MECH-SM-FRAC is identical to MECH-SM when the original problem is used as a relaxation (and then, of course, $\rho = 1$).

MECH-SM-FRAC$(A, v, \mathbf{c}, B)$

Set $A' = \{i \mid c_i \leq B\}$ and $i^* \in \arg\max_{i \in A'} v(i)$
if $\left(\rho + 1 + \sqrt{\rho^2 + 4\rho + 1}\right) \cdot v(i^*) \geq \text{OPT}_f(A \setminus \{i^*\}, v, \mathbf{c}_{-i^*}, B)$ **then**
    return i^*
else
    return GREEDY-SM$(A, v, \mathbf{c}, B/2)$

A.5 Proving Theorems 2.5 and 2.6

Monotonicity (and therefore truthfulness and individual rationality) and budget-feasibility of the three mechanisms directly follow from [9] and [3]. What is left to show are the approximation guarantees.

The key fact here is the following lemma from [18]. Note that the lemma also follows from Lemma D.2 for $\epsilon = 0$ and $\beta = 0.5$. For the rest of this subsection, S will denote the outcome of GREEDY-SM$(A, v, \mathbf{c}, B/2)$.

Lemma A.5. $\text{OPT}(A, v, \mathbf{c}, B) \leq 3 \cdot v(S) + 2 \cdot v(i^*)$.

For RAND-MECH-SM, if X denotes the outcome of the mechanism, then directly by Lemma A.5 we have $\text{E}(X) \geq \frac{3}{5} v(S) + \frac{2}{5} v(i^*) \geq \frac{1}{5} \text{OPT}(A, v, \mathbf{c}, B)$, thus proving the approximation ratio.

For MECH-SM we consider two cases:

- If i^* is returned by the mechanism, then $(2+\sqrt{6}) \cdot v(i^*) \geq \text{OPT}(A \setminus \{i^*\}, v, \mathbf{c}_{-i^*}, B) \geq \text{OPT}(A, v, \mathbf{c}, B) - v(i^*)$, and therefore $\text{OPT}(A, v, \mathbf{c}, B) \leq (3+\sqrt{6}) \cdot v(i^*)$.

- On the other hand, if S is returned, then $(2+\sqrt{6})\cdot v(i^*) < \text{OPT}(A\backslash\{i^*\}, v, \mathbf{c}_{-i^*}, B) \leq \text{OPT}(A, v, \mathbf{c}, B)$. Combining this with Lemma A.5 we have $\text{OPT}(A, v, \mathbf{c}, B) \leq 3 \cdot v(S) + \frac{2}{2+\sqrt{6}} \text{OPT}(A, v, \mathbf{c}, B)$ and therefore $\text{OPT}(A, v, \mathbf{c}, B) \leq \frac{3(2+\sqrt{6})}{\sqrt{6}} \cdot v(S) = (3+\sqrt{6}) \cdot v(S)$.

For MECH-SM-FRAC we also consider two cases. Let $\alpha = \rho + 1 + \sqrt{\rho^2 + 4\rho + 1}$.

- If i^* is returned, then $\alpha \cdot v(i^*) \geq \text{OPT}_f(A\backslash\{i^*\}, v, \mathbf{c}_{-i^*}, B) \geq \text{OPT}(A\backslash\{i^*\}, v, \mathbf{c}_{-i^*}, B) \geq \text{OPT}(A, v, \mathbf{c}, B) - v(i^*)$, and therefore $\text{OPT}(A, v, \mathbf{c}, B) \leq (\alpha+1) \cdot v(i^*)$.
- On the other hand, if S is returned, then $\alpha \cdot v(i^*) < \text{OPT}_f(A \setminus \{i^*\}, v, \mathbf{c}_{-i^*}, B) \leq \rho \cdot \text{OPT}(A \setminus \{i^*\}, v, \mathbf{c}_{-i^*}, B) \leq \rho \cdot \text{OPT}(A, v, \mathbf{c}, B)$. Combining this with Lemma A.5 we have $\text{OPT}(A, v, \mathbf{c}, B) \leq 3 \cdot v(S) + \frac{2\rho}{\alpha} \text{OPT}(A, v, \mathbf{c}, B)$ and therefore $\text{OPT}(A, v, \mathbf{c}, B) \leq \frac{3\alpha}{\alpha - 2\rho} \cdot v(S) = (\alpha+1) \cdot v(S)$, where the last equality is only a matter of calculations.

A.6 Regarding the Closing Remark of Section 2

We use the cut function on a very simple graph and show that although v is submodular, $\hat{v}$ is not. Consider the following graph where each edge has unit weight:

We compute the value of the following sets:

$\hat{v}(\{a\}) = v(\{a\}) = 2$

$\hat{v}(\{a,b\}) = v(\{a\}) = 2$

$\hat{v}(\{a,c\}) = v(\{a\}) = v(\{c\}) = v(\{a,c\}) = 2$

$\hat{v}(\{a,b,c\}) = v(\{b,c\}) = 3$

Now it is easy to see that although $\{a\} \subseteq \{a,c\}$ we have $\hat{v}(\{a\} \cup \{b\}) - \hat{v}(\{a\}) = 0 < 1 = \hat{v}(\{a,c\} \cup \{b\}) - \hat{v}(\{a,c\})$.

So an interesting question is if $\hat{v}$ can be classified when v is submodular. In [15] the general XOS class was introduced, where it is allowed for a function to be non-monotone (recall here that the XOS class contains only non-decreasing functions by definition). They proved that when v is general XOS then $\hat{v}$ is XOS, while in addition they observed that the class of non-negative submodular functions is a strict subset of the general XOS class. Thus since any non-negative submodular function v is also general XOS we conclude $\hat{v}$ is XOS.

B Missing Proofs from Section 3

Before proving Lemma 3.1 and Theorem 3.2, we begin with a simple fact and a useful lemma.

Fact B.1. *For any* $S \subseteq A$, $\max\{\text{OPT}(S,B), \text{OPT}(A\backslash S,B)\} \geq 0.5\,\text{OPT}(A,B)$ *since* $\text{OPT}(S,B)+\text{OPT}(A\backslash S,B) \geq \text{OPT}(A,B)$ *by subadditivity.*

Lemma B.2. *For any* $S \subseteq A$, $\text{OPT}(S,\infty) \leq 2n \cdot \text{OPT}(A,B)$.

Proof. Recall that $|\{i \in A \mid c_i > B\}| \leq 1$. Let $S^* \subseteq S$ be such that $v(S^*) = \text{OPT}(S,\infty)$. By subadditivity we have $\text{OPT}(S,\infty) = v(S^*) \leq \sum_{i\in S^*} v(i)$. Consider three cases.

If $\{i \in A \mid c_i > B\} = \emptyset$, then by the fact that every singleton is a feasible solution we have $\sum_{i\in S^*} v(i) \leq n \cdot \max_{i\in A} v(i) \leq n \cdot \text{OPT}(A,B)$.

If $\{i \in A \mid c_i > B\} = \{x\} \nsubseteq S^*$, then every singleton in $A\backslash\{x\}$ is a feasible solution, and like before we have $\sum_{i\in S^*} v(i) \leq (n-1) \cdot \max_{i\in A\backslash\{x\}} v(i) \leq (n-1) \cdot \text{OPT}(A,B)$.

If $\{i \in A \mid c_i > B\} = \{x\} \subseteq S^*$, then we need to bound $v(x)$. Since v is symmetric we have $v(x) = v(A\backslash\{x\}) \leq \sum_{i\in A\backslash\{x\}} v(i) \leq (n-1) \cdot \max_{i\in A\backslash\{x\}} v(i)$. Therefore, by using again that every singleton in $A\backslash\{x\}$ is a feasible solution, we have $\sum_{i\in S^*} v(i) \leq v(x) + (n-1) \cdot \max_{i\in A\backslash\{x\}} v(i) \leq 2n \cdot \max_{i\in A\backslash\{x\}} v(i) \leq 2n \cdot \text{OPT}(A,B)$. □

Proof of Lemma 3.1. By Fact B.1 we have $\text{OPT}(X,B) \geq 0.5\,\text{OPT}(A,B)$. Let $T \subseteq X\backslash\{i\}$ for some $i \in X$. By submodularity we have $v(T\cup\{i\}) - v(T) \geq v(X) - v(X\backslash\{i\})$. Since S is a $(1+\frac{\epsilon}{4n^2})$-approximate local optimum and v is symmetric, X is also a $(1+\frac{\epsilon}{4n^2})$-approximate local optimum. As a result, $v(X\backslash\{i\}) \leq (1+\frac{\epsilon}{4n^2})v(X)$ and thus $v(X) - v(X\backslash\{i\}) \geq -\frac{\epsilon}{4n^2}v(X) \geq -\frac{\epsilon}{4n^2}\text{OPT}(X,\infty) \geq -\frac{\epsilon}{2n}\text{OPT}(A,B) \geq -\frac{\epsilon}{n}\text{OPT}(X,B)$, where the second to last inequality follows from Lemma B.2. □

Proof of Theorem 3.2. To facilitate the exposition of the analysis, we restate algorithm GREEDY-ENUM-SM, so that instead of the set variable T, we have a variable S^t that describes the set constructed at iteration t.

VAR-GREEDY-ENUM-SM$(A, v, \mathbf{c}, B)$

Let S_1 be the best feasible solution of cardinality at most 3 (by enumerating them all)
$S_2 = \emptyset$
for *every* $U \subseteq A$ *with* $|U| = 3$ **do**
  $S^0 = U,\ t = 1,\ A^0 = A\backslash U$
  while $A^{t-1} \neq \emptyset$ **do**
    Find $\theta_t = \max_{i\in A^{t-1}} \dfrac{v(S^{t-1}\cup\{i\}) - v(S^{t-1})}{c_i}$, and let i_t be an element of A^{t-1} that attains θ_t
    if $\theta_t \geq 0$ *and* $\sum_{i\in S^{t-1}\cup\{i_t\}} c_i \leq B$ **then**
      $S^t = S^{t-1}\cup\{i_t\}$
    else
      $S^t = S^{t-1}$
    $A^t = A^{t-1}\backslash\{i_t\}$
    $t = t+1$
  if $v(S^{t-1}) > v(S_2)$ **then**
    $S_2 = S^{t-1}$
return $S \in \arg\max_{X\in\{S_1,S_2\}} v(X)$

Recall that VAR-GREEDY-ENUM-SM runs on both S and $A\backslash S$ and LS-GREEDY returns the best solution of these two. We may assume, without loss of generality that $\text{OPT}(S,B) = \max\{\text{OPT}(S,B), \text{OPT}(A\backslash S,B)\}$ (the case for $A\setminus S$ being symmetric). By Fact B.1 we have $\text{OPT}(S,B) \geq 0.5\,\text{OPT}(A,B)$. So, it suffices to show that running VAR-GREEDY-ENUM-SM on S outputs a set of value at least $(1-1/e-\epsilon)\,\text{OPT}(S,B)$.

In what follows we analyze the approximation ratio achieved by VAR-GREEDY-ENUM-SM$(S,v,\mathbf{c},B)$ with respect to $\text{OPT}(S,B)$. For this, we follow closely the proof of the main result in [29].

If there is an optimal solution for our problem restricted on S, of cardinality one, two or three, then the set S_1 of VAR-GREEDY-ENUM-SM will be such a solution. Hence, assume that the cardinality of any optimal solution is at least four and let S^* be such a solution. If necessary, reorder the elements of $S^* = \{j_1,\ldots,j_{|S^*|}\}$ so that $j_1 = \arg\max_\ell v(\{j_\ell\})$, and $j_{k+1} = \arg\max_{\ell>k}\left[v(\{j_1,\ldots,j_k,j_\ell\}) - v(\{j_1,\ldots,j_k\})\right]$

Let $Y = \{j_1, j_2, j_3\}$. For notational convenience, we will use the function $g(\cdot) = v(\cdot) - v(Y)$. It is straightforward that $g(\cdot)$ is submodular. Moreover, the following fact follows from [29].

Fact B.3. *$g(X\cup\{i\}) - g(X) \leq \frac{1}{3}v(Y)$, for any $Y\subseteq X\subseteq S$ and $i\in S^*\setminus X$.*

Consider the execution of the greedy algorithm with initial set $U = Y$. Let t^*+1 be the first time when an element $i_{t^*+1}\in S^*$ is not added to S^{t^*}. In fact, we assume that t^*+1 is the first time when $S^t = S^{t-1}$. (To see that this is without loss of generality, if there is some time $\tau < t^*+1$ such that i_τ is not added to $S^{\tau-1}$, then—by the definition of t^*+1—it must be the case that $i_\tau\notin S^*$. But then, we may consider the instance $(S\setminus\{i_\tau\}, v, \mathbf{c}_{S\setminus\{i_\tau\}}, B)$ instead. We have $v(S^*) = \text{OPT}(S\setminus\{i_\tau\},B) = \text{OPT}(S,B)$ and the greedy solution constructed in the iteration where $S_0 = Y$ is exactly the same as before.) We are going to distinguish two cases.

Case 1. For all $t\in[t^*], \theta_t\geq 0$, but $\theta_{t^*+1} < 0$. Using Theorem 2.2 for S^* and S^{t^*} we have

$$\begin{aligned}
g(S^*) &\leq g(S^{t^*}) + \sum_{i\in S^*\setminus S^{t^*}} (g(S^{t^*}\cup\{i\}) - g(S^{t^*})) - \sum_{i\in S^{t^*}\setminus S^*} (g(S^{t^*}\cup S^*) - g(S^{t^*}\cup S^*\setminus\{i\})) \\
&= g(S^{t^*}) + \sum_{i\in S^*\setminus S^{t^*}} (v(S^{t^*}\cup\{i\}) - v(S^{t^*})) - \sum_{i\in S^{t^*}\setminus S^*} (v(S^{t^*}\cup S^*) - v(S^{t^*}\cup S^*\setminus\{i\})) \\
&\leq g(S^{t^*}) + \sum_{i\in S^*\setminus S^{t^*}} c_i\theta_{t^*+1} - |S^{t^*}\setminus S^*|\left(-\frac{\epsilon}{n}\,\text{OPT}(S,B)\right) \\
&\leq g(S^{t^*}) + \epsilon\,\text{OPT}(S,B)\,,
\end{aligned}$$

Here, the second to last inequality holds by Lemma 3.1 and by the assumptions we have made. That is, for every $i\in S^*\setminus S^{t^*}$, we have that $i\in A^{t^*}$, since we assumed that t^*+1 is the first time when $S^t = S^{t-1}$, hence up until time t^*, A^{t^*} contains all the agents apart from S^{t^*}. This implies that for every $i\in S^*\setminus S^{t^*}$, we have that $v(S^{t^*}\cup\{i\}) - v(S^{t^*}) \leq c_i\theta_{t^*+1}$, by the definition of θ_{t^*+1}.

Therefore, we can conclude that

$$v(S^{t^*}) = v(Y) + g(S^{t^*}) \geq v(Y) + g(S^*) - \epsilon\,\text{OPT}(S,B) = (1-\epsilon)\,\text{OPT}(S,B)\,.$$

Case 2. For all $t\in[t^*+1], \theta_t\geq 0$, but $\sum_{i\in S^{t^*}\cup\{i_{t^*+1}\}} c_i > B$ while $\sum_{i\in S^{t^*}} c_i \leq B$. Using Theorem 2.2 for S^* and each of S^t, $t\in[t^*]$, as well as Lemma 3.1, we have

$$\begin{aligned}
g(S^*) &\leq g(S^t) + \sum_{i\in S^*\setminus S^t} (g(S^t\cup\{i\}) - g(S^t)) - \sum_{i\in S^t\setminus S^*} (g(S^t\cup S^*) - g(S^t\cup S^*\setminus\{i\})) \\
&= g(S^t) + \sum_{i\in S^*\setminus S^t} (v(S^t\cup\{i\}) - v(S^t)) - \sum_{i\in S^t\setminus S^*} (v(S^t\cup S^*) - v(S^t\cup S^*\setminus\{i\}))
\end{aligned}$$

$$\leq g(S^t) + \sum_{i \in S^* \setminus S^t} (v(S^t \cup \{i\}) - v(S^t)) - |S^t \setminus S^*| \left(-\frac{\epsilon}{n} \operatorname{OPT}(S,B)\right)$$
$$\leq g(S^t) + \sum_{i \in S^* \setminus S^t} (v(S^t \cup \{i\}) - v(S^t)) + \epsilon \operatorname{OPT}(S,B),$$

and therefore

$$\begin{aligned} g(S^*) - \epsilon \operatorname{OPT}(S,B) &\leq g(S^t) + \sum_{i \in S^* \setminus S^t} (v(S^t \cup \{i\}) - v(S^t)) \\ &\leq g(S^t) + \sum_{i \in S^* \setminus S^t} c_i \theta_{t+1} \\ &\leq g(S^t) + \left(B - \sum_{i \in Y} c_i\right) \theta_{t+1}, \end{aligned}$$

for all $t \in [t^*]$.

For the last part of the proof we need the following inequality of Wolsey [30].

Theorem B.4 (Wolsey [30])**.** *Let k and s be arbitrary positive integers, and $\rho_1, \ldots, \rho_k$ be arbitrary reals with $z_1 = \sum_{i=1}^{k} \rho_i$ and $z_2 = \min_{t \in [k]} \left(\sum_{i=1}^{t-1} \rho_i + s\rho_t\right) > 0$. Then $z_1/z_2 \geq 1 - (1-1/s)^k \geq 1 - e^{-\frac{k}{s}}$.*

For any τ, define $B_\tau = \sum_{t=1}^{\tau} c_{i_t}$, and let $k = B_{t^*+1}$ and $s = B - \sum_{i \in Y} c_i$. We also define $\rho_1, \ldots, \rho_k$ as follows; for $i \leq c_1$, $\rho_i = \theta_1$ and for $B_\tau < i \leq B_{\tau+1}$, $\rho_i = \theta_{\tau+1}$. It is easy to see that $g(S^{t^*} \cup \{i_{t^*+1}\}) = \sum_{t=1}^{t^*+1} c_{i_t} \theta_{i_t} = \sum_{i=1}^{k} \rho_i$ and similarly $g(S^t) = \sum_{t=1}^{\tau} c_{i_t} \theta_{i_t} = \sum_{i=1}^{B_\tau} \rho_i$.

As noted in [29], since the ρ_is are nonnegative we have

$$\min_{t \in [k]} \left(\sum_{i=1}^{t-1} \rho_i + s\rho_t \right) = \min_{\tau \in [t^*]} \left(\sum_{i=1}^{B_\tau} \rho_i + s\rho_{B_\tau+1} \right),$$

and therefore

$$g(S^*) - \epsilon \operatorname{OPT}(S,B) \leq \min_{t \in [k]} \left(\sum_{i=1}^{t-1} \rho_i + s\rho_t \right).$$

So, as a direct application of Theorem B.4 we have

$$\frac{g(S^{t^*} \cup \{i_{t^*+1}\})}{g(S^*) - \epsilon \operatorname{OPT}(S,B)} \geq 1 - e^{-\frac{k}{s}} > 1 - e^{-1}.$$

Finally, using the above inequality and Fact B.3, we get

$$\begin{aligned} v(S^{t^*}) &= v(Y) + g(S^{t^*}) = v(Y) + g(S^{t^*} \cup \{i_{t^*+1}\}) - (g(S^{t^*} \cup \{i_{t^*+1}\}) - g(S^{t^*})) \\ &\geq v(Y) + (1-e^{-1}) g(S^*) - (1-e^{-1}) \epsilon \operatorname{OPT}(S,B) - (v(S^{t^*} \cup \{i_{t^*+1}\}) - v(S^{t^*})) \\ &\geq v(Y) + (1-e^{-1}) g(S^*) - \epsilon \operatorname{OPT}(S,B) - \frac{1}{3} v(Y) \\ &\geq (1 - e^{-1} - \epsilon) \operatorname{OPT}(S,B). \end{aligned}$$

Since in both cases the final output T^* of the algorithm has value at least $v(S^{t^*})$, this implies that

$$v(T^*) \geq (1 - e^{-1} - \epsilon) \operatorname{OPT}(S,B) \geq \frac{1 - e^{-1} - \epsilon}{2} \operatorname{OPT}(A,B) > \left(\frac{e-1}{2e} - \epsilon\right) \operatorname{OPT}(A,B),$$

thus concluding the analysis of the performance of the algorithm. □

C Missing Proofs from Section 4

Proof of Lemma 4.1. The fact that v is submodular when restricted on 2^S is trivial. Suppose now that the statement is not true and that v is not non-decreasing on 2^S. That is, there exist $T, T' \subseteq S$ such that $T \subsetneq T'$ and $v(T) > v(T')$. Let $T' \setminus T = \{i_1, \ldots, i_r\}$.

By Theorem 2.2 we have

$$v(T) \leq v(T') - \sum_{j \in [r]} (v(T') - v(T' \setminus \{i_j\})),$$

and therefore

$$\sum_{j \in [r]} (v(T') - v(T' \setminus \{i_j\})) \leq v(T') - v(T) < 0.$$

We conclude that there is some $\ell \in [r]$ such that $v(T') - v(T' \setminus \{i_\ell\}) < 0$. Then, by submodularity and the fact that $T' \setminus \{i_\ell\} \subseteq S \setminus \{i_\ell\}$, we get

$$v(S) - v(S \setminus \{i_\ell\}) \leq v(T') - v(T' \setminus \{i_\ell\}) < 0.$$

But then $v(S \setminus \{i_\ell\}) > v(S)$, which contradicts the fact that S is a local maximum of v. So, it must be the case that v is non-decreasing on the subsets of S. □

Proof of Theorem 4.2. The fact that both mechanism RAND-MECH-SYMSM and DET-MECH-SYMSM are budget-feasible follows from the budget-feasibility of RAND-MECH-SM and MECH-SM respectively, established in [9].

For truthfulness and individual rationality, it suffices to show that the allocation rule is monotone. Let us look first at DET-MECH-SYMSM. Consider an agent i with true cost c_i in an instance I where i is included in the winning set. Note first that the local search step is not affected by the costs, hence no player can influence the local optimum. Suppose now that it was the case that $\text{OPT}(S, B) \geq \text{OPT}(A \setminus S, B)$, hence $i \in S$. If player i now declares a lower cost, then the optimal solution within S can only get better, hence the mechanism will run MECH-SM on S as before. Since MECH-SM is monotone, player i will again be selected in the solution. We conclude that the outcome rule is monotone and DET-MECH-SYMSM is truthful.

To prove that the randomized mechanism RAND-MECH-SYMSM is universally truthful we use similar arguments. We fix the random bits of the mechanism and we consider a winning agent i like before. Again, no player can influence the outcome of local search. Suppose it is the case that $\text{OPT}(S, B) \geq \text{OPT}(A \setminus S, B)$, hence $i \in S$. If i declares a lower cost, then the optimal solution within S improves, and RAND-MECH-SM will still run on S. Since RAND-MECH-SM is universally truthful, it is monotone given the random bits, and player i will again be a winner. We conclude that RAND-MECH-SYMSM is universally truthful.

To argue now about the approximation ratio, suppose that we are in the case that $\text{OPT}(S, B) \geq \text{OPT}(A \setminus S, B)$ (the other case being symmetric). We know by Fact B.1 that $\text{OPT}(S, B) \geq 0.5 \cdot \text{OPT}(A, B)$. Hence, since we run either MECH-SM or RAND-MECH-SM on S, we will get twice their approximation ratio. This implies a ratio of 10 for RAND-MECH-SYMSM and a ratio of $6 + 2\sqrt{6}$ for DET-MECH-SYMSM. □

Proof of Fact 4.3. The lower bounds of [9] are both for additive objectives (Knapsack). It suffices to show that given an instance $(A, v, \mathbf{c}, B)$ of Knapsack, we can construct an equivalent $(A', v', \mathbf{c}', B')$ instance of Budgeted Max Weighted Cut. This is straightforward. Consider a graph G with vertex set $A \cup x$ and edge set $\{(i, x) \mid i \in A\}$. For $i \in A$, vertex i has cost $c'_i = c_i$, while vertex x has cost $c'_x = B + 1$. Edge (i, x) has weight $v(i)$. Finally, for $S \subseteq A \cup x$, $v'(S)$ is equal to the weight of the cut defined by S.

This correspondence between items and vertices creates a natural correspondence between solutions. It is clear that each feasible solution of the Knapsack instance essentially defines a feasible solution to the Budgeted Max Weighted Cut instance of the same value and vice versa. In particular $\text{OPT}(A, v, \mathbf{c}, B) = \text{OPT}(A', v', \mathbf{c}', B)$. We conclude that any lower bound for Knapsack gives a lower bound for Budgeted Max Weighted Cut. □

Proof of Theorem 4.4. For completeness, mechanism RAND-MECH-UCUT is stated below. For the unweighted version of Max Cut, the mechanism runs in polynomial time. The fact that the mechanism is truthful, individually rational and budget feasible follows from the same arguments as in the proof of Theorem 4.2.

RAND-MECH-UCUT$(A, v, \mathbf{c}, B)$

1 $S =$ APPROX-LOCAL-SEARCH$(A, v, 0)$ //find an exact local optimum
2 with probability 1/2 **return** RAND-MECH-SM$(S, v, \mathbf{c}_S, B)$
3 with probability 1/2 **return** RAND-MECH-SM$(A \setminus S, v, \mathbf{c}_{A\setminus S}, B)$

Finally, let $X_S = \text{GREEDY-SM}(S, v, \mathbf{c}_S, B/2)$, $X_{A\setminus S} = \text{GREEDY-SM}(A \setminus S, v, \mathbf{c}_{A\setminus S}, B/2)$, $i_S \in \arg\max_{\{i \in S \mid c_i \le B\}} v(i)$, and $i_{A\setminus S} \in \arg\max_{\{i \in A\setminus S \mid c_i \le B\}} v(i)$.

Since we run mechanism RAND-MECH-SM with probability 1/2 on S and $A \setminus S$, it is not hard to see that each of X_S and $X_{A\setminus S}$ is returned with probability 3/10, while each of i_S and $i_{A\setminus S}$ is returned with probability 2/10. If X denotes the outcome of the mechanism, then using subadditivity and Lemma A.5, we get

$$\begin{aligned} \mathrm{E}(X) &= \frac{3}{10} v(X_S) + \frac{2}{10} v(i_S) + \frac{3}{10} v(X_{A\setminus S}) + \frac{2}{10} v(i_{A\setminus S}) \\ &\ge \frac{1}{10} \text{OPT}(S, v, \mathbf{c}_S, B) + \frac{1}{10} \text{OPT}(A \setminus S, v, \mathbf{c}_{A\setminus S}, B) \\ &\ge \frac{1}{10} \text{OPT}(A, v, \mathbf{c}, B)\,, \end{aligned}$$

thus establishing an approximation ratio of 10. □

Proof of Theorem 4.5. The proof takes the same approach as the proof of Lemma 3 from [3]. We begin with a nonlinear program such that if all the x_is are integral then the objectives (1) and (8) have the same value.

$$\text{maximize: } F(x) = \sum_{i \in [n]} \sum_{j \in [n] \setminus [i]} w_{ij}(x_i + x_j - 2x_i x_j) \tag{8}$$

$$\text{subject to: } \sum_{i \in [n]} c_i x_i \le B \tag{9}$$

$$x_i = 0\,, \quad \forall i \in [n] \setminus X \tag{10}$$

$$0 \le x_i \le 1\,, \quad \forall i \in [n] \tag{11}$$

So, if x is any feasible integral vector to our problem, we have $F(x) \leq \text{OPT}(I)$. Moreover, given any feasible solution x, z to (1)-(6), the value of (1) is upper bounded by $L(x) = \sum_{i\in[n]} \sum_{j\in[n]\setminus[i]} w_{ij} \min\{x_i + x_j, 2 - x_i - x_j\}$, since $z_{ij} \leq \min\{x_i + x_j, 2 - x_i - x_j\}$ for any $i \in [n]$, $j \in [n] \setminus [i]$.

Next, we show that $F(x) \geq 0.5L(x)$. This follows from the inequality

$$2(a + b - 2ab) \geq \min\{a + b, 2 - a - b\} \text{ for any } a, b \in [0, 1],$$

proven by Ageev and Sviridenko [1]. For completeness we prove it here as well. Notice that by replacing a and b by $1 - a$ and $1 - b$ respectively both sides of the inequality remain exactly the same. Therefore, it suffices to prove $2(a + b - 2ab) \geq a + b$ for any $a, b \in [0, 1]$ such that $a + b \leq 1$ (since otherwise, $1 - a + 1 - b \leq 1$). This however is equivalent to $a + b \geq 4ab$ which is true since $a + b \geq (a + b)^2 \geq 4ab$.

Hence, if x^*, z^* is an optimal fractional solution to (1)-(6), then the value of (1) is $\text{OPT}_f(I) = L(x^*)$ and thus $F(x^*) \geq 0.5L(x^*) = 0.5\,\text{OPT}_f(I)$. However, x^* may have several fractional coordinates. Our next step is to transform x^* to a vector x' that has at most one fractional coordinate and at the same time $F(x') \geq F(x^*)$. To this end, we show how to reduce the fractional coordinates by (at least) one in any feasible vector with at least two such coordinates.

Consider a feasible vector x, and suppose x_i and x_j are two non integral coordinates. Note that $i, j \in X$. Let $x_\epsilon^{i,j}$ be the vector we get if we replace x_i by $x_i + \epsilon$ and x_j by $x_j - \epsilon c_i / c_j$ and leave every other coordinate of x the same. Note that the function $\bar{F}(\epsilon) = F(x_\epsilon^{i,j})$, with respect to ϵ, is either linear or a polynomial of degree 2 with positive leading coefficient. That is, $\bar{F}(\epsilon)$ is convex.

Notice now that $x_\epsilon^{i,j}$ always satisfies the budget constraint (9), and also satisfies (11) as long as $\epsilon \in \left[\max\left\{-x_i, (x_j - 1)c_j / c_i\right\}, \min\left\{1 - x_i, x_j c_j / c_i\right\}\right]$. Due to convexity, $\bar{F}(\epsilon)$ attains a maximum on one of the endpoints of this interval, say at ϵ^*. Moreover, at either endpoint at least one of $x_i + \epsilon^*$ and $x_j - \epsilon^* c_i / c_j$ is integral. That is, $x_{\epsilon^*}^{i,j}$ has at least one more integral coordinate than x and $F(x_{\epsilon^*}^{i,j}) \geq F(x)$.

So, initially $x \leftarrow x^*$. As long as there exist two non integral coordinates x_i and x_j we set $x \leftarrow x_{\epsilon^*}^{i,j}$ as described above. This happens at most $n - 1$ times, and outputs a feasible vector x' with at most one non-integral coordinate, and with $F(x') \geq F(x^*)$. We have then the following implications:

$$\text{OPT}_f(I) = L(x^*) \leq 2 \cdot F(x^*) \leq 2 \cdot F(x'). \tag{12}$$

If x' is integral, then by (12) we have $\text{OPT}_f(I) \leq 2 \cdot F(x') \leq 2 \cdot \text{OPT}(I)$, and we are done. So, suppose that x'_r is the only fractional coordinate of x'. Let x^0 and x^1 be the vectors we get if we set x'_r to 0 or 1 respectively and leave every other coordinate of x' the same. Notice that for $a, b \in [0, 1]$ the inequality $(1 - a)(b - 1) \leq 0$ implies $a + b - ab \leq 1$ and therefore $a + b - 2ab \leq 1$, so we have

$$F(x') - F(x^0) = \sum_{j\in[n]\setminus r} w_{rj}(x'_r - 2x'_r x'_j) \leq \sum_{j\in[n]\setminus r} w_{rj}(x'_r + x'_j - 2x'_r x'_j) \leq \sum_{j\in[n]\setminus r} w_{rj} = F(x^1 - x^0),$$

and thus

$$F(x') \leq F(x^0) + F(x^1 - x^0). \tag{13}$$

Combining (12)and (13), we have

$$\text{OPT}_f(I) \leq 2 \cdot F(x') \leq 2 \cdot \left(F(x^0) + F(x^1 - x^0)\right),$$

Using the fact that F is upper bounded by $\text{OPT}(I)$ on integral vectors, we have that $F(x^0) \leq \text{OPT}(I)$. Observe now also that $F(x^1 - x^0) = \sum_{j \in [n] \setminus r} w_{rj} = v(r) \leq \beta_I \, \text{OPT}(I)$, by the definition of β_I. Hence, overall we get

$$\text{OPT}_f(I) \leq (2 + 2\beta_I) \, \text{OPT}(I) \,,$$

thus completing the proof. □

Proof of Theorem 4.6. Clearly the mechanism runs in polynomial time.

For truthfulness and individual rationality, it suffices to show that the allocation rule is monotone, i.e. a winning agent j remains a winner if he decreases his cost to $c'_j < c_j$. If $j = i^*$ and he wins in line 2, then his bid is irrelevant and he remains a winner. If $j \neq i^*$, or $j = i^*$ but he wins after line 3, we may assume he wins at line 7 (the case of line 9 is symmetric). When bidding $c'_j < c_j$, the choice of the mechanism in line 3 does not change (if $j \neq i^*$ then OPT_f is improved, if $j = i^*$ nothing changes). Further, since he cannot influence the local search, j is still in S and MECH-SM-FRAC is executed. By the monotonicity of MECH-SM-FRAC we have that j is still a winner. Therefore, the mechanism is monotone.

If the winner is i^* in line 3, then his payment is B. Otherwise, budget-feasibility follows from the budget-feasibility of MECH-SM-FRAC and the observation that the comparison in line 3 only gives additional upper bounds on the payments of winners from MECH-SM-FRAC.

It remains to prove the approximation ratio. We consider two cases. Let $\alpha = 26.25$. If i^* is returned in line 3, then

$$\alpha \cdot v(i^*) \geq \text{OPT}_f(A \setminus \{i^*\}, B) \geq \text{OPT}(A \setminus \{i^*\}, B) \geq \text{OPT}(A, B) - v(i^*) \,,$$

and therefore $\text{OPT}(A, B) \leq (\alpha + 1) \cdot v(i^*) = 27.25 \cdot v(i^*)$.

On the other hand, if $X = \text{MECH-SM-FRAC}(S, B)$ is returned in line 7, then by Theorem 4.5 with factor 4 we have

$$\alpha \cdot v(i^*) < \text{OPT}_f(A \setminus \{i^*\}, B) \leq 4 \cdot \text{OPT}(A \setminus \{i^*\}, B) \leq 4 \cdot \text{OPT}(A, B).$$

Therefore, $v(i^*) < \frac{4}{\alpha} \, \text{OPT}(A, B)$ and for the remaining steps of the mechanism we can use Theorem 4.5 with factor $2 + 8/\alpha$.

At line 6 it must be the case $\text{OPT}_f(S, B) \geq \text{OPT}_f(A \setminus S, B)$. Thus,

$$\left(2 + \frac{8}{\alpha}\right) \text{OPT}(S, B) \geq \text{OPT}_f(S, B) \geq \text{OPT}_f(A \setminus S, B) \geq \text{OPT}(A \setminus S, B) \geq \text{OPT}(A, B) - \text{OPT}(S, B).$$

Therefore $\text{OPT}(S, B) \geq \frac{\alpha}{3\alpha + 8} \, \text{OPT}(A, B)$. By Theorem 2.6 we have

$$\left(4 + 8/\alpha + \sqrt{(2 + 8/\alpha)^2 + 32/\alpha + 9}\right) v(X) \geq \text{OPT}(S, B) \geq \frac{\alpha}{3\alpha + 8} \, \text{OPT}(A, B) \,.$$

and by substituting α and doing the calculations we get $\text{OPT}(A, B) \leq 27.25 \cdot v(X)$.

The case where $X = \text{MECH-SM-FRAC}(A \setminus S, B)$ is returned in line 9 is symmetric to the above and we do not need to consider it separately. We conclude that in any case $\text{OPT}(A, B) \leq 27.25 \cdot$ DET-MECH-UCUT(A, B). □

D Missing Material from Section 5

D.1 Going Beyond Non-Decreasing Objectives

Here we pave the way for the main results of Section 5, and therefore we need to prove that certain mechanisms work even for objectives that are not exactly non-decreasing. To make this precise, given a ground set A, a budget B and a constant $\epsilon \geq 0$, we say that a set function v is (B,ϵ)-*quasi-monotone* (or just quasi-monotone) on a set $X \subseteq A$ if for every $T \subsetneq X$ and every $i \in X \setminus T$, we have $v(T \cup \{i\}) - v(T) \geq -\frac{\epsilon}{n} \textsc{opt}(X, v, \mathbf{c}, B)$. Clearly, $(B,0)$-quasi-monotone on X just means non-decreasing on X.

The main lemmas needed for our proofs are about Greedy-SM, as it all boils down to the monotonicity, budget-feasibility, and approximation ratio of this simple mechanism. As mentioned in Lemma A.4, Greedy-SM is monotone, since any item out of the winning set remains out of the winning set if it increases its cost.

Lemma D.1. *Suppose v is a (B,ϵ)-quasi-monotone submodular function on A such that $U \cdot v(S) \geq \textsc{opt}(A, v, \mathbf{c}, B)$, where S is the set output by* Greedy-SM$(A, v, \mathbf{c}, (1-U\epsilon)B/2)$ *and U is a constant.*[4] *Assuming the payments of Myerson's lemma, S is budget-feasible.*

Proof. Recall that before the description of Greedy-SM we assumed—without loss of generality—that agents are sorted in descending order with respect to their ratio of marginal value over cost, i.e. $1 = \arg\max_{j \in A} \frac{v(j)}{c_j}$ and $i = \arg\max_{j \in A \setminus [i-1]} \frac{v([j]) - v([j-1])}{c_j}$ for $i \geq 2$. Suppose that $S = [\ell]$, i.e., $1, 2, \ldots, \ell$ are added in S in that order. Let $S_0 = \emptyset = [0]$ and $S_i = [i]$ for $1 \leq i \leq \ell$. We are going to show that the payment to agent i is upper bounded by $\frac{B \cdot (v([i]) - v([i-1]))}{v(S)}$, and then budget feasibility directly follows from $\sum_{i \in [\ell]} \frac{B \cdot (v([i]) - v([i-1]))}{v(S)} = B$.

Suppose this upper bound does not hold for every agent. That is, there exists some $j \in [\ell]$ such that agent j bids $b_j > \frac{B \cdot (v([j]) - v([j-1]))}{v(S)}$ and is still included in the output S' of Greedy-SM$(A, v, (\mathbf{c}_{-j}, b_j), (1-U\epsilon)B/2)$. Let $\mathbf{b} = (\mathbf{c}_{-j}, b_j)$ and notice that up to agent $j-1$, agents are added to S' in the same order as they do in S, but after that the ordering might be affected. Also for $i \in [\ell]$ we have $\frac{v([i]) - v([i-1])}{c_i} \geq \frac{v([\ell]) - v([\ell-1])}{c_\ell} \geq \frac{2 \cdot v([\ell])}{(1-U\epsilon)B}$ and therefore

$$\mathbf{b}(S \setminus \{j\}) = \mathbf{c}(S \setminus \{j\}) \leq \mathbf{c}(S) = \sum_{i \in [\ell]} c_i \leq \frac{(1-U\epsilon)B}{2 \cdot v([\ell])} \sum_{i \in [\ell]} (v([i]) - v([i-1])) = \frac{(1-U\epsilon)B}{2}.$$

For Greedy-SM$(A, v, (\mathbf{c}_{-j}, b_j), (1-U\epsilon)B/2)$ let S'_{j-1} be the chosen set right before j is added and $S'_j = S'_{j-1} \cup \{j\}$. This implies

$$j \in \arg\max_{i \in [n]} \frac{v(S'_{j-1} \cup \{i\}) - v(S'_{j-1})}{b_i} \text{ and } \frac{v(S'_j) - v(S'_{j-1})}{b_j} \geq \frac{2 \cdot v(S'_j)}{(1-U\epsilon)B}.$$

By Theorem 2.2 we have

$$v(S) - v(S'_j) \leq \sum_{i \in S \setminus S'_j} (v(S'_j \cup \{i\}) - v(S'_j)) - \sum_{i \in S'_j \setminus S} (v(S'_j \cup S) - v(S'_j \cup S \setminus \{i\}))$$

[4] Typically, in our case, U is a constant associated with the constant ρ that determines how the optimal solution to the relaxed problem is bounded by the optimal solution to the original problem. In particular, throughout this work, U is upper bounded by $(1+\rho)(2+\rho+\sqrt{\rho^2+4\rho+1})+1$.

$$\leq \sum_{i\in S\setminus S'_j} (v(S'_j\cup\{i\})-v(S'_j)) - |S'_j\setminus S|\left(-\frac{\epsilon}{n}\,\mathrm{OPT}(A,v,\mathbf{c},B)\right)$$
$$\leq \sum_{i\in S\setminus S'_j} (v(S'_j\cup\{i\})-v(S'_j)) + \epsilon U\cdot v(S)\,.$$

If $S\setminus S'_j=\emptyset$ then we directly get $v(S'_j)\geq(1-\epsilon U)\,v(S)$, otherwise we have

$$\begin{aligned}
(1-\epsilon U)\,v(S)-v(S'_j) &\leq \sum_{i\in S\setminus S'_j} b_i\cdot\frac{v(S'_j\cup\{i\})-v(S'_j)}{b_i}\\
&\leq \max_{i\in S\setminus S'_j}\frac{v(S'_j\cup\{i\})-v(S'_j)}{b_i}\cdot\sum_{i\in S\setminus S'_j} b_i\\
&\leq \max_{i\in[n]}\frac{v(S'_j\cup\{i\})-v(S'_j)}{b_i}\cdot\mathbf{b}(S\setminus S'_j)\\
&\leq \max_{i\in[n]}\frac{v(S'_{j-1}\cup\{i\})-v(S'_{j-1})}{b_i}\cdot\mathbf{b}(S\setminus\{j\})\\
&\leq \frac{v(S'_{j-1}\cup\{j\})-v(S'_{j-1})}{b_j}\cdot\frac{(1-U\epsilon)B}{2}\\
&\leq \frac{v(S_{j-1}\cup\{j\})-v(S_{j-1})}{b_j}\cdot\frac{(1-U\epsilon)B}{2}\\
&\leq \frac{v(S)}{B}\cdot\frac{(1-U\epsilon)B}{2}=\frac{(1-U\epsilon)\,v(S)}{2}\,.
\end{aligned}$$

Therefore, $v(S'_j)\geq(1/2-\epsilon U+\epsilon U/2)\,v(S)$.

In any case, we have $v(S'_j)\geq(1/2-U\epsilon/2)\,v(S)$ and thus

$$\frac{v(S_j)-v(S_{j-1})}{b_j}\geq\frac{v(S'_j)-v(S'_{j-1})}{b_j}\geq\frac{2v(S'_j)}{(1-U\epsilon)B}\geq\frac{2(1/2-U\epsilon/2)\,v(S)}{(1-U\epsilon)B}\geq\frac{v(S)}{B}\,,$$

which contradicts our assumption about b_j. □

Lemma D.2. *Suppose v is a (B,ϵ)-quasi-monotone submodular function on A. Let S be the set output by* GREEDY-SM$(A,v,\mathbf{c},\beta B)$ *and* $i^*\in\operatorname{argmax}_{i\in\{j\in A\,|\,c_j\leq B\}} v(i)$. *Then* $\mathrm{OPT}(A,v,\mathbf{c},B)\leq\frac{1}{1-\epsilon}\left(\frac{1+\beta}{\beta}v(S)+\frac{1}{\beta}v(i^*)\right)$.

Proof. Like in the proof of Lemma D.1 we assume that agents are sorted in descending order with respect to their ratio of marginal value over cost, and that $S=[\ell]$. Let S^* be an optimal budget-feasible solution, i.e., $v(S^*)=\mathrm{OPT}(A,v,\mathbf{c},B)$. By Theorem 2.2 we have

$$\begin{aligned}
v(S^*)-v(S) &\leq \sum_{i\in S^*\setminus S}(v(S\cup\{i\})-v(S)) - \sum_{i\in S\setminus S^*}(v(S\cup S^*)-v(S\cup S^*\setminus\{i\}))\\
&\leq \sum_{i\in S^*\setminus S} c_i\cdot\frac{v(S\cup\{i\})-v(S)}{c_i} - |S\setminus S^*|\left(-\frac{\epsilon}{n}\,\mathrm{OPT}(A,v,\mathbf{c},B)\right)\\
&\leq \max_{i\in S^*\setminus S}\frac{v(S\cup\{i\})-v(S)}{c_i}\cdot\sum_{i\in S^*\setminus S} c_i+\epsilon v(S^*)
\end{aligned}$$

$$
\begin{aligned}
&= \frac{v(S\cup\{\ell+1\})-v(S)}{c_{\ell+1}}\cdot \mathbf{c}(S^*\setminus S)+\epsilon v(S^*)\\
&\le \frac{v(S\cup\{\ell+1\})}{\beta B}\cdot \mathbf{c}(S^*)+\epsilon v(S^*) \le \frac{v(S\cup\{\ell+1\})}{\beta}+\epsilon v(S^*)\\
&\le \frac{1}{\beta}(v(S)+v(\ell+1))+\epsilon v(S^*) \le \frac{1}{\beta}(v(S)+v(i^*))+\epsilon v(S^*)\,.
\end{aligned}
$$

By rearranging the terms we get $\textsc{opt}(A,v,\mathbf{c},B) \le \frac{1}{1-\epsilon}\left(\frac{1+\beta}{\beta}v(S)+\frac{1}{\beta}v(i^*)\right)$. □

We close with an analog of Theorem for a variant of MECH-SM-FRAC where instead of

D.2 Proofs of Theorems 5.1 and 5.3

Proof of Theorem 5.1. We are going to need a slight variant of MECH-SM-FRAC:

MECH-SM-FRAC-VAR$(A,v,\mathbf{c},B,\gamma)$

1 Set $A'=\{i \mid c_i \le B\}$ and $i^* \in \arg\max_{i\in A'} v(i)$
2 **if** $\left(\rho+1+\sqrt{\rho^2+4\rho+1}\right)\cdot v(i^*) \ge \textsc{opt}_f(A\setminus\{i^*\},v,\mathbf{c}_{-i^*},B)$ **then**
3 **return** i^*
4 **else**
5 **return** GREEDY-SM$(A,v,\mathbf{c},\gamma B/2)$

Recall that $\alpha=(1+\rho)\left(2+\rho+\sqrt{\rho^2+4\rho+1}\right)-1$. Now we are ready to state mechanism DET-MECH-SYMSM-FRAC. Clearly the mechanism runs in polynomial time.

DET-MECH-SYMSM-FRAC$(A,v,\mathbf{c},B)$

1 Set $A'=\{i \mid c_i \le B\}$ and $i^* \in \arg\max_{i\in A'} v(i)$
2 **if** $\alpha\cdot v(i^*) \ge \textsc{opt}_f(A\setminus\{i^*\},v,\mathbf{c}_{-i^*},B)$ **then**
3 **return** i^*
4 **else**
5 $S=$ APPROX-LOCAL-SEARCH(A,v,ϵ')
6 **if** $\textsc{opt}_f(S,v,\mathbf{c}_S,B) \ge \textsc{opt}_f(A\setminus S,v,\mathbf{c}_{A\setminus S},B)$ **then**
7 **return** MECH-SM-FRAC-VAR$(S,v,\mathbf{c}_S,B,(1-(\alpha+2)\epsilon'))$
8 **else**
9 **return** MECH-SM-FRAC-VAR$(A\setminus S,v,\mathbf{c}_{A\setminus S},B,(1-(\alpha+2)\epsilon'))$

The proof follows the proof of Theorem 4.6. In fact, the monotonicity—and thus truthfulness and individual rationality—of the mechanism follows from that proof and the observation that MECH-SM-FRAC-VAR is monotone even when v is non-monotone. The latter is due to the monotonicity of GREEDY-SM which is straightforward and briefly discussed in Appendix D.1.

We proceed with the approximation ratio. If i^* is returned in line 3, then

$$
\alpha\cdot v(i^*) \ge \textsc{opt}_f(A\setminus\{i^*\},B) \ge \textsc{opt}(A\setminus\{i^*\},B) \ge \textsc{opt}(A,B)-v(i^*)\,,
$$

and therefore $\textsc{opt}(A,B) \le (\alpha+1)\cdot v(i^*)$.

On the other hand, if X is returned by MECH-SM-FRAC-VAR in line 7, then we have

$$\alpha \cdot v(i^*) < \text{OPT}_f(A \setminus \{i^*\}, B) \leq \rho \cdot \text{OPT}(A \setminus \{i^*\}, B) \leq \rho \cdot \text{OPT}(A, B).$$

Therefore, $v(i^*) < \frac{\rho}{\alpha} \text{OPT}(A, B)$.

At line 6 it must be the case $\text{OPT}_f(S, B) \geq \text{OPT}_f(A \setminus S, B)$. Thus,

$$\rho \cdot \text{OPT}(S, B) \geq \text{OPT}_f(S, B) \geq \text{OPT}_f(A \setminus S, B) \geq \text{OPT}(A \setminus S, B) \geq \text{OPT}(A, B) - \text{OPT}(S, B).$$

Therefore $\text{OPT}(S, B) \geq \frac{1}{\rho+1} \text{OPT}(A, B)$. Now we need the following lemma about the performance of MECH-SM-FRAC-VAR.

Lemma D.3. *For $\eta = \rho + 1 + \sqrt{\rho^2 + 4\rho + 1}$ and any $\epsilon'' > 0$, there is a sufficiently small ϵ' so that $\text{OPT}(S, B) \leq (\eta + 1 + \epsilon'')$MECH-SM-FRAC-VAR$(S, v, \mathbf{c}_S, B, (1 - (\alpha + 2)\epsilon'))$. Moreover ϵ' only depends on the constants ρ and ϵ''.*

Proof. Let $\delta = (1 - (\alpha + 2)\epsilon')/2$. For MECH-SM-FRAC-VAR we also consider two cases.

If i^* is returned, then $\eta \cdot v(i^*) \geq \text{OPT}_f(A \setminus \{i^*\}, v, \mathbf{c}_{-i^*}, B) \geq \text{OPT}(A \setminus \{i^*\}, v, \mathbf{c}_{-i^*}, B) \geq \text{OPT}(A, v, \mathbf{c}, B) - v(i^*)$, and therefore $\text{OPT}(A, v, \mathbf{c}, B) \leq (\eta + 1) \cdot v(i^*)$.

On the other hand, if the outcome X of GREEDY-SM is returned, then $\eta \cdot v(i^*) < \text{OPT}_f(A \setminus \{i^*\}, v, \mathbf{c}_{-i^*}, B) \leq \rho \cdot \text{OPT}(A \setminus \{i^*\}, v, \mathbf{c}_{-i^*}, B) \leq \rho \cdot \text{OPT}(A, v, \mathbf{c}, B)$. Combining this with Lemma D.2 we have $\text{OPT}(A, v, \mathbf{c}, B) \leq \frac{1}{1-\epsilon'} \left(\frac{1+\delta}{\delta} v(X) + \frac{1}{\delta} \frac{\rho}{\eta} \text{OPT}(A, v, \mathbf{c}, B) \right)$ and therefore, for sufficiently small ϵ', $\text{OPT}(A, v, \mathbf{c}, B) \leq \left(\frac{3\eta}{\eta - 2\rho} + \epsilon'' \right) v(S) = (\eta + 1 + \epsilon'') \cdot v(S)$, where the calculations are the same as in the proof of the approximation ratio of MECH-SM-FRAC. □

Now, combining all the above, we have

$$\text{OPT}(A, B) \leq (\rho + 1) \text{OPT}(S, B) \leq (\rho + 1) \left(\rho + 2 + \epsilon'' + \sqrt{\rho^2 + 4\rho + 1} \right) v(X).$$

For $\epsilon'' = \frac{\epsilon}{\rho+1}$ we have $(\rho + 1) \left(\rho + 2 + \epsilon'' + \sqrt{\rho^2 + 4\rho + 1} \right) = \alpha + 1 + \epsilon$, as desired. So, the ϵ' used in the mechanism is the one we get from the proof of Lemma D.3 for $\epsilon'' = \frac{\epsilon}{\rho+1}$.

The case where X is returned in line 9 is symmetric. We conclude that in any case $\text{OPT}(A, B) \leq (\alpha + 1 + \epsilon) \cdot$ DET-MECH-SYMSM-FRAC(A, B).

It remains to show that the mechanism is budget feasible. If the winner is i^* in line 3, then his payment is B. Otherwise, we would like budget-feasibility to follow from the budget-feasibility of MECH-SM-FRAC-VAR and the observation that the comparison in line 3 only gives additional upper bounds on the payments of winners from MECH-SM-FRAC. However, the budget-feasibility of MECH-SM-FRAC-VAR depends on the budget-feasibility of GREEDY-SM, and according to Lemma D.1 we need $v(X) \geq \frac{1}{\alpha+2} \cdot \text{OPT}(A, v, \mathbf{c}, B)$, where X is the output of GREEDY-SM. This, however, follows from the approximation ratio for $\epsilon \leq 1$. Thus the mechanism is budget-feasible. □

Proof of Theorem 5.3. The proof is identical with the proof of Theorem 5.1 with the exception of the analysis of the approximation ratio which borrows the ideas of the proof of Theorem 4.6. So we only focus on the approximation ratio.

If i^* is returned in line 3, then $\text{OPT}(A, B) \leq (\alpha_1 + 1) \cdot v(i^*)$ like before.

Otherwise, using Theorem 4.5, we have

$$\alpha_1 \cdot v(i^*) < \text{OPT}_f(A\setminus\{i^*\}, B) \leq 4 \cdot \text{OPT}(A\setminus\{i^*\}, B) \leq 4 \cdot \text{OPT}(A, B).$$

Therefore, $v(i^*) < \frac{4}{\alpha_1}\,\text{OPT}(A,B)$ and for the remaining steps of the mechanism we can use Theorem 4.5 with factor $\rho' = 2 + 8/\alpha_1$.

Without loss if generality, assume that X is returned by MECH-SM-FRAC-VAR in line 7. At line 6 it must be the case $\text{OPT}_f(S,B) \geq \text{OPT}_f(A\setminus S, B)$. Thus, following the previous analysis, $\text{OPT}(S,B) \geq \frac{1}{\rho'+1}\,\text{OPT}(A,B)$.

Combining this inequality with Lemma D.3 for ρ' we get that for $\epsilon'' = \frac{\epsilon}{\rho'+1}$ there is an ϵ' such that

$$\text{OPT}(A,B) \leq (\rho'+1)\left(\rho' + 2 + \epsilon'' + \sqrt{\rho'^2 + 4\rho' + 1}\right) v(X) = (\alpha_1 + 1 + \epsilon) \cdot v(X)\,.$$

This ϵ' is to be used in the mechanism. Showing that $\alpha_1 = 26.245$ works is only a matter of calculations, and for $\epsilon \leq 1/200$ we get a 27.25-approximate solution. □

E Missing Material from Section 6

Proof of Theorem 6.2. Our mechanism has the same structure as the one presented in [6] but we tune its parameters and perform a slightly different analysis in order to improve the approximation upper bound. Initially we revisit the *random sampling* part and modify the threshold bound of line 3:

SAMPLE-XOS

1 Pick each item indepedently at random with probability $\frac{1}{2}$ into group T
2 Compute an otpimal solution $\text{OPT}(T)$ for items in T given budget B
3 Set a threshold $t = \frac{v(\text{OPT}(T))}{4{,}64B}$
4 Consider items in $A\setminus T$ and find a set $S^* \in \arg\max_{S\subseteq A\setminus T}\{v(S) - t\cdot c(S)\}$
5 Find additive function α with $\alpha(S^*) = v(S^*)$ in the XOS representation of $v(\cdot)$
6 Run ADDITIVE-MECHANISM for function $\alpha(\cdot)$ with respect to set S^* and budget B
7 Output the result of ADDITIVE-MECHANISM

This part is used as one of the two alternatives of the main mechanism. We modify the probabilities with which the two outcomes occur:

MAIN-XOS

1 With probability 0.916, run SAMPLE-XOS
2 With probability 0.084, pick a most valuable item as the only winner and pay him B

Universal truthfulness, individual rationality and budget feasibility follow directly from the proof of Theorem 3.1 of [6].

We proceed in proving the approximation ratio of the mechanism. Let $i^* \in \arg\max_{i\in A} v(i)$, X_S be the output of SAMPLE-XOS, $\text{OPT} = \text{OPT}(A,B)$ and $\text{OPT}^* = \text{OPT}_\alpha(S^*,B)$, i.e., the optimal solution with respect to the item set S^*, additive function $\alpha(\cdot)$ and budget B.

We will also need the following:

Claim E.1 (Claim 3.1 of [6])**.** *For any $S \subseteq S^*$, $\alpha(S) - t\cdot c(S) \geq 0$.*

Lemma E.2 (Lemma 2.1 of [6])**.** *Consider any subadditive function $v(\cdot)$. For any given subset $S \subseteq A$ and a positive integer k assume that $v(S) \geq k \cdot v(i)$ for any $i \in S$. Further, suppose that S is divided uniformly at random into two groups T_1 and T_2. Then with probability at least $\frac{1}{2}$, we have that $v(T_1) \geq \frac{k-1}{4k} v(S)$ and $v(T_2) \geq \frac{k-1}{4k} v(S)$.*

For the rest of the proof we set $k = 29$ and thus we get $v(\text{OPT}(T)) \geq v(\text{OPT} \cap T) \geq \frac{28}{4 \cdot 29} v(\text{OPT}) = \frac{7}{29} v(\text{OPT})$ with a probability of at least $\frac{1}{2}$.

We follow a similar analysis as [6] but since we set $k = 29$, things get a llitle more complicated. Let X_M to be the output of MAIN-XOS. We have the following three case regarding the value of item i^*:

1. $v(i^*) > \frac{2}{29} v(\text{OPT})$. In this case we have that $v(X_M) = 0.084 \cdot v(i^*) + 0.916 \cdot v(X_S) > 0.084 \cdot v(i^*) = 0.084 \cdot \frac{2}{29} v(\text{OPT}) = 0.0058 \cdot v(\text{OPT}) \Rightarrow v(\text{OPT}) \leq 172,41 \cdot v(X_M)$.

2. $\frac{1}{29} v(\text{OPT}) < v(i^*) \leq \frac{2}{29} v(\text{OPT})$. We split this case into two subcases:

   - $C(S^*) > B$.
     Since $C(S^*)$ is higher than B we can find a subset $S' \subseteq S^*$, such that $\frac{B}{2} \leq c(S') \leq B$. In addition by using Claim 3.1 of [6] we can derive that $\alpha(S') \geq t \cdot c(S') \geq \frac{v(\text{OPT}(T)) \cdot B}{2 \cdot 4.64 \cdot B} \geq \frac{v(\text{OPT}(T))}{9.28}$. So since OPT^* is an optimal solution and S' a particular solution under budget constraint B, we can conclude that $\alpha(\text{OPT}^*) \geq \alpha(S') > \frac{v(\text{OPT}(T))}{9.28} \geq \frac{7 \cdot v(\text{OPT})}{29 \cdot 9.28} = 0.026 \cdot v(\text{OPT})$ with probability at least $\frac{1}{2}$.
   - $C(S^*) \leq B$.
     In this case we have that $\text{OPT}^* = S^* = \arg\max_{S \subseteq A \setminus T} \{v(S) - t \cdot c(S)\}$. Now let $S' = \text{OPT} \setminus T$, so $c(S') \leq c(\text{OPT}) \leq B$ and thus by using Lemma 2.1 of [6] we get that $v(S') \geq \frac{7}{29} v(\text{OPT})$ with probability of at least $\frac{1}{2}$. Then, we have that: $\alpha(\text{OPT}^*) = \alpha(S^*) = v(S^*) \geq v(S^*) - t \cdot c(S^*) \geq v(S') - t \cdot c(S') \geq \frac{7}{29} v(\text{OPT}) - \frac{v(\text{OPT}(T)}{4.64 \cdot B} \cdot B \geq \frac{7}{29} v(\text{OPT}) - \frac{1}{4.64} \cdot v(\text{OPT}) = 0.0259 \cdot v(\text{OPT})$ with probability at least $\frac{1}{2}$.

   In both cases we have that $\alpha(\text{OPT}^*) \geq 0.0259 \cdot v(\text{OPT})$ with probability at least $\frac{1}{2}$. Now recall that ADDITIVE-MECHANISM has an approximation factor of at most 3 with respect to the optimal solution $\alpha(\text{OPT}^*)$. So we can finally derive that $v(X_S) \geq \frac{1}{3} \alpha(\text{OPT}^*) \geq \frac{1}{3} \cdot \frac{1}{2} \cdot 0.0259 \cdot v(\text{OPT}) = 0.0044 \cdot v(\text{OPT})$. Thus the solution that MAIN-XOS has value $v(X_M) = 0.084 \cdot v(i^*) + 0.916 \cdot v(X_S) \geq \frac{1}{29} \cdot 0.084 \cdot v(\text{OPT}) + 0.916 \cdot 0.0044 \cdot v(\text{OPT}) = 0,007 \cdot v(\text{OPT})$ and we can get $v(\text{OPT}) \leq 142,86 \cdot v(X_M)$.

3. $v(i^*) \leq \frac{1}{29} v(\text{OPT})$. The analysis of case 2 holds here as well, we only have to revisit the final part since now we have a different bound for $v(i^*)$. Specifically, $v(X_M) = 0.084 \cdot v(i^*) + 0.916 \cdot v(X_S) \geq 0.916 \cdot v(X_S) \geq 0.916 \cdot 0.0044 \cdot v(\text{OPT}) = 0,004 \cdot v(\text{OPT}) \Rightarrow v(\text{OPT}) \leq 250 \cdot v(X_M)$

So by examining all the possible cases, we can conclude that $v(\text{OPT}) \leq 250 \cdot v(X_M)$. □